

VERITAS NetBackup Vault™ 4.5

System Administrator's Guide

for UNIX and Windows

March 2002
30-000526-011


VERITAS

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

Copyright

Copyright © 1993-2002 VERITAS Software Corporation. All Rights Reserved. VERITAS, VERITAS SOFTWARE, the VERITAS logo, *Business Without Interruption*, VERITAS The Data Availability Company, and VERITAS NetBackup Vault are trademarks or registered trademarks of VERITAS Software Corporation in the U.S. and/or other countries. Other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
Phone 650-527-8000
Fax 650-527-8050
www.veritas.com



Contents

Preface	xi
Organization	xi
Related Documents	xii
Online Documentation	xiii
On Windows NT/2000	xiii
On UNIX	xiii
On the Support Web Site	xiv
Conventions	xiv
Type Style	xiv
Notes and Cautions	xiv
Key Combinations	xv
Command Usage	xv
Terms	xvi
Getting Help	xvi
Accessibility	xvii
 Chapter 1. Introduction	1
Why Use Vault?	1
How Vault Meets these Requirements	2
What Vault Does	2
How Vault Fits into the NetBackup Product	3
How to Access Vault	3
What is Vaulting?	4
Vaulting Duplicate or Original Media	4



What Is the Primary Copy?	5
What is Duplication?	5
What is a Catalog Backup?	6
What is the Eject Step?	6
What is the Reporting Step?	6
Terminology	7
Chapter 2. Installing NetBackup Vault 4.5	9
Supported Platforms	9
Supported Robots	9
Installation Prerequisites	10
Installing on UNIX Systems	10
Installing on Windows Systems	11
Upgrading	11
Uninstalling	12
On UNIX Systems	12
On Windows Systems	12
Chapter 3. Preparing to Use Vault	15
Notes and Recommendations	15
Managing Multiple Media Types within a Vault	15
Collecting NetBackup Information	16
Vault Policy Types	17
Gathering NetBackup Information	18
Robot Information	20
Media Manager Configuration	20
Network Configuration for Receiving Reports	20
Vault Configuration	21
Chapter 4. Best Practices	25
Preferred Strategies	26



Preferred Strategy: Vault Original Backups	26
Disk Staging	26
Make Sure You Vault All You Want To	27
Overlap the Time Window in the Profile	27
Resolve Multiple Names for a Single Server	28
Specify Robotic Volume Group When Configuring Logical Vault	28
Ensuring Report Integrity	29
Further Delineating Reports	30
Making a Report Match Media Assigned to a Logical Vault	30
Do Not Vault More Than You Want To	31
Send Only the Intended Backups Off-site	31
Use the Suspend Option to Avoid Vaulting Partial Backups	32
Vaulting Original Backups in a 24x7 Environment	33
Preparing for Efficient Recovery	33
Revault Media	33
Use Good Naming Conventions for Volume Pools	34
Specify a Unique Volume Pool For Each Logical Vault	34
Specifying the Primary Copy and Keeping It On Site	34
Avoid Resource Contention During Duplication	34
When Two Processes Try to Use the Same Drive	35
When the Read Drive Is Not in the Vault's Robot	37
Sharing Resources with Backup Jobs	38
Load Balancing	39
Specifying Different Volume Pools for Source and Destination	40
Avoid Sending Duplicates Over The Network	40
Use ITC (Multiple Copies) To Minimize Data Transfer	40
Use Alternate Read Server	41
Use Advanced Duplication Configuration	41
Increase Duplication Throughput	43
Configuring for Multiple-Drives: Basics	43



Multiple-Drive Scenario: Does Not Send Data Over Network	43
Chapter 5. Configuring Vault	45
Methods of Configuration	45
When to Use vltadm	45
Configuring Robots for Vault	46
Creating a Vault	47
Creating a Vault Policy	48
Creating a Profile	50
Configuring a Profile	51
Configuring the Duplication Tab	54
Configuring the Catalog Backup Tab	64
Backing up the Catalog Using Vault	66
Configuring the Eject Tab	68
Configuring the Reports Tab	69
Chapter 6. Inline Tape Copy	73
Understanding Multiple Copies (Inline Tape Copy)	73
Selecting Continue vs. Fail for Multiple Copies	74
Using Inline Tape Copy (Multiple Copies) in Vault	76
Creating Multiple Copies from Outside Vault	79
Configuring Inline Tape Copy through the Policy Node	79
Configuring Inline Tape Copy Through Catalog Node	81
Chapter 7. Administrative Tasks in Vault	87
What Is a Vault Session?	87
Running a Vault Session	88
Resuming a Vault Session	91
Entering Addresses for Email Notification	92
Editing a Vault or Profile	93
Printing Vault and Profile Information	93



Copying a Profile	94
Moving a Vault to a Different Robot	94
Adding Alternate Media Server Names	95
Understanding Log Files	96
Output from the Vault Driver	98
Setting the Duration of Vault Working Files	99
Setting the Duration and Level of Vault Log Files	99
Using Notify Scripts	100
Notify Script for a Specific Profile	102
Notify Script for a Specific Vault	102
Notify Script for a Specific Robot	102
Order of Execution	103
Vault Support in Activity Monitor	103
Extended Error Codes	105
Ensuring Available Media for Catalog Backups	106
Manual Deassigning of Vaulted NetBackup Catalog Media	107
Chapter 8. Reporting	109
Report Types	109
Reports for Media Going Off Site	110
Reports for Media Coming On Site	111
Detailed Media Reports	111
Recovery Report for Vault	113
Consolidating Reports	113
Running Reports from the Command Line	114
Report Distribution	115
Reprinting Reports	115
Chapter 9. Using the Menu User Interfaces (MUIs)	117
Menu User Interfaces in Vault 4.5	117
Accessing the Menu User Interfaces	118



Using vltadm	118
Help for vltadm	128
How to Use the Help Screens	128
Using vltopmenu	129
Changes in vmadm for Vault	131
Additions to Volume Configuration	131
Changes to the Special Actions Menu	131
Changes to bpdbjobs for Vault	134
Chapter 10. Troubleshooting	135
Errors Returned by the Vault Session	135
Other Troubleshooting Issues	135
Media Missing in Robot	135
Bad or Missing Duplicate Tape	136
If You Need to Stop Vault	136
Tape Drive or Robot Offline	137
No Duplicate Progress Message	137
Ejecting Tapes While in Use	138
Ejecting More Media Than Export Capacity	138
Vault Session Locking	138
Logs To Accompany Problem Reports	139
Appendix A. Commands	141
User Requirements	141
vltadm	142
vlteject	144
vltinject	148
vltoffsitemedias	150
vltopmenu	153
vltrun	154



Appendix B. Upgrading bpvault to NetBackup Vault 4.5	159
Introduction to Vault 4.5	159
If You Choose Not to Upgrade	160
Feature Comparison between Versions	160
Feature Descriptions	163
Where to Find Data Between Versions	166
Comparison of Variables and Actions in Each Version	166
Comparison of Parameters in Each Version	168
Cross References: Executables, Directories, and Files	171
Appendix C. Recovering Damaged Media	173
Steps in Recovering Backup Images	173
Following the Image Recovery Procedure	174
Revaulting Media After a Restore	178
Appendix D. Vault's File and Directory Structure	181
Directories and Files Created During Installation	181
Appendix E. Functional Design Appendix for Vault	187
Functional Design Overview	187
Overview	187
Other Related Services	188
Other Related Vault Documents	188
Architectural Services	188
Services Interactions Diagram	190
Client/Server Architectural Services	191
Architectural Example	193
Technical Components	194
Components for Vault	194
Technical Design Issues	194
Vault Technical Components	195



Technical Example: Standard Duplication Diagram	203
Technical Example: Standard Duplication Table	204
Operational Procedures	211
NetBackup Vault Image Duplication Process	214
Glossary	217
Index	245



Preface

The VERITAS mission for NetBackup 4.5 is to provide heterogeneous data protection solutions from the workgroup or departmental level to the enterprise level. NetBackup BusinessServer 4.5 is a high-performance solution offering ease of use and functionality that is appropriate to smaller work environments. NetBackup DataCenter 4.5 provides powerful functionality, flexibility, mainframe caliber robustness, and high performance in large enterprise environments. NetBackup 4.5 is a major new release in the NetBackup product family with significant enhancements in the area of scalability, application protection, disaster recovery, server-free and off-host backup, and integration with Backup Exec.

NetBackupVault was created to simplify the processes of image duplication, offsite storage, and offsite retrieval for both storage administrators and systems operators. This *System Administrator's Guide* details the responsibilities of and explains procedures performed by the system administrator. It provides the administrative information needed to run VERITAS NetBackup Vault on both UNIX and Windows platforms. A separate *Operator's Guide* provides detailed instructions for system operations.

Organization

This guide is organized as follows:

- ◆ Chapter 1, "Introduction," introduces the Vault product, providing detailed discussions about background, design and the latest features.
- ◆ Chapter 2, "Installing NetBackup Vault 4.5," presents the steps required to install and configure Vault.
- ◆ Chapter 3, "Preparing to Use Vault," provides a pre-planning checklist to complete in preparation for installing and configuring Vault.
- ◆ Chapter 4, "Best Practices," provides advice on how to configure Vault efficiently.
- ◆ Chapter 5, "Configuring Vault," describes procedures for configuring Vault robots, vaults, and profiles.
- ◆ Chapter 6, "Inline Tape Copy," presents detailed information on using the Inline Tape Copy feature, which allows you to configure multiple simultaneous backup images.



- ◆ Chapter 7, “Administrative Tasks in Vault,” covers the various tasks involved in using Vault.
- ◆ Chapter 8, “Reporting,” details the reports available through Vault, how they are generated, and how to receive notification of Vault activity.
- ◆ Chapter 9, “Using the Menu User Interfaces (MUIs),” explains the Vault functionality available through the two menu-driven interfaces provided.
- ◆ Chapter 10, “Troubleshooting,” discusses potential problems that may occur when using Vault and how to resolve or work around them.
- ◆ Appendix A, “Commands,” describes commands available through Vault.
- ◆ Appendix B, “Upgrading bpvault to NetBackup Vault 4.5,” compares and contrasts bpvault 3.4 and NetBackup Vault 4.5.
- ◆ Appendix C, “Recovering Damaged Media,” describes the steps to take to recovery images from damaged or missing media.
- ◆ Appendix D, “Vault’s File and Directory Structure,” describes the directories and files installed with the Vault product.
- ◆ Appendix E, “Functional Design Appendix for Vault,” describes the architectural and technical components of Vault.

Related Documents

The following documents provide related information. For a more detailed listing of NetBackup documents, refer to *NetBackup Release Notes*.

If you have a UNIX server, refer to these documents:

- ◆ *NetBackup Release Notes*
Provides important information about NetBackup software, such as the platforms and operating systems that are supported and operating notes that may not be in the manuals or the online help.
- ◆ *NetBackup DataCenter System Administrator’s Guide - UNIX*
Explains how to configure and manage NetBackup DataCenter on a UNIX system.
- ◆ *NetBackup Troubleshooting Guide - UNIX*
Provides troubleshooting information for UNIX-based NetBackup products.

If you have a Windows server, refer to these documents:

- ◆ *NetBackup Release Notes*

Provides important information about NetBackup software, such as the platforms and operating systems that are supported and operating notes that may not be in the manuals or the online help.

◆ *NetBackup DataCenter System Administrator's Guide - Windows NT/2000*

Explains how to configure and manage NetBackup DataCenter on a Windows server system.

◆ *NetBackup Troubleshooting Guide - Windows*

Provides troubleshooting information for Windows-based NetBackup products.

Online Documentation

On Windows NT/2000

The released software contains on-line PDF and ASCII versions of these release notes and a readme file for the client. If you choose to install the documentation during setup, NetBackup installs these documents in the following locations on your disk:

◆ *install_path\Help*

Adobe Acrobat Portable Document Format (PDF) copies of all related documents, including these release notes.

◆ The readme files on *install_path\NetBackup* are:

- *Readme.txt* (The *Readme.txt* file (ASCII format) may be slightly more up-to-date than the printed and pdf copies of the release notes.)
- *Readme_Client.txt*
- *Readme_Server.txt*
- *Readme_SMS.txt*
- *Readme_Win2000.txt*

On UNIX

The product CD-ROM also contains PDF copies of these release notes and other documents.

Note You will need Adobe Acrobat Reader to view the PDF documents. The latest version of Acrobat Reader is available on the Adobe web site:
<http://www.adobe.com>.
VERITAS assumes no responsibility for the correct installation or use of the reader.



On the Support Web Site

Copies of NetBackup documentation are also available on the VERITAS support web site:

1. Go to the VERITAS support web page
`www.support.veritas.com/`
2. In the VERITAS Support Product List, choose NetBackup Products.
3. A page appears with a list of the NetBackup products. Choose NetBackup BusinessServer or NetBackup DataCenter.
4. The documents page appears. Choose the document you want.

Conventions

The following explains typographical and other conventions used in this guide.

Type Style

Typographic Conventions

Typeface	Usage
Bold fixed width	Input. For example, type cd to change directories.
Fixed width	Paths, commands, filenames, or output. For example: The default installation directory is <code>/opt/VRTSxx</code> .
<i>Italics</i>	Book titles, new terms, or used for emphasis. For example: <i>Do not</i> ignore cautions.
<i>Sans serif</i> (italics)	Placeholder text or variables. For example: Replace <i>filename</i> with the name of your file.
Serif (no italics)	Graphical user interface (GUI) objects, such as fields, menu choices, etc. For example: Enter your password in the Password field.

Notes and Cautions

Note This is a Note. Notes are used to call attention to information that makes using the product easier or helps in avoiding problems.

Caution This is a Caution. Cautions are used to warn about situations that could cause data loss.

Key Combinations

Some keyboard command sequences use two or more keys at the same time. For example, holding down the **Ctrl** key while pressing another key. Keyboard command sequences are indicated by connecting the keys with a plus sign. For example:

Press Ctrl+t

Command Usage

The following conventions are frequently used in the synopsis of command usage.

brackets []

The enclosed command line component is optional.

Vertical bar or pipe (|)

Separates optional arguments from which the user can choose. For example, when a command has the following format:

`command arg1 | arg2`

the user can use either the *arg1* or *arg2* variable.



Terms

The terms listed in the table below are used in the VERITAS NetBackup documentation to increase readability while maintaining technical accuracy.

Term	Definition
Microsoft Windows, Windows	<p>Terms used as nouns to describe a line of operating systems developed by Microsoft, Inc.</p> <p>A term used as an adjective to describe a specific product or noun. Some examples are: Windows 95, Windows 98, Windows NT, Windows 2000, Windows servers, Windows clients, Windows platforms, Windows hosts, and Windows GUI.</p> <p>Where a specific Windows product is identified, then only that particular product is valid with regards to the instance in which it is being used.</p> <p>For more information on the Windows operating systems that NetBackup supports, refer to the VERITAS support web site at http://www.support.veritas.com.</p>
Windows servers	<p>A term that defines the Windows server platforms that NetBackup supports; those platforms are: Windows NT and Windows 2000.</p>
Windows clients	<p>A term that defines the Windows client platforms that NetBackup supports; those platforms are: Windows 95, 98, ME, NT, 2000, XP (for 32- and 64-bit versions), and LE.</p>

Getting Help

For updated information about this product, including system requirements, supported platforms, supported peripherals, and a list of current patches available from Technical Support, visit our web site:

<http://www.support.veritas.com/>



VERITAS Customer Support has an extensive technical support structure that enables you to contact technical support teams that are trained to answer questions to specific products. You can contact Customer Support by sending an e-mail to support@veritas.com, or by finding a product-specific phone number from the VERITAS support web site. The following steps describe how to locate the proper phone number.

1. Open <http://www.support.veritas.com/> in your web browser.
2. Click **Contact Support**. The *Contacting Support Product List* page appears.
3. Select a product line and then a product from the lists that appear. The page will refresh with a list of technical support phone numbers that are specific to the product you just selected.

Accessibility

NetBackup contains features that make the user interface easier to use by people who are vision impaired and by people who have limited dexterity. Accessibility features include:

- ◆ Support for assistive technologies such as screen readers and voice input (Windows servers only)
- ◆ Support for keyboard (mouse-less) navigation using accelerator keys and mnemonic keys

For more information about accessibility in NetBackup, see the NetBackup system administrator's guide.





This chapter introduces NetBackup Vault, an off-site backup management extension to NetBackup. Topics covered include:

- ◆ Why Use Vault?
- ◆ How Vault Meets these Requirements
- ◆ What Vault Does
- ◆ How Vault Fits into the NetBackup Product
- ◆ How to Access Vault
- ◆ What is Vaulting?
- ◆ Vaulting Duplicate or Original Media
- ◆ What Is the Primary Copy?
- ◆ What is Duplication?
- ◆ What is a Catalog Backup?
- ◆ What is the Eject Step?
- ◆ What is the Reporting Step?
- ◆ Terminology

Why Use Vault?

One of the first steps in a disaster recovery plan is to establish an off-site backup management system. In the context of NetBackup Vault, a vault is an off-site storage location that holds critical data. This data is normally stored on magnetic media (tapes) and replicates data located at the data center.

If a disaster such as a flood, fire, tornado, or hurricane were to wipe out an entire site, these systems could be rebuilt on new hardware if and only if a copy of the mission-critical data were stored safely off site.

Other requirements of an off-site backup management strategy are:



- ◆ A schedule for sending media off site
- ◆ A way to determine which media to retrieve from off-site storage for reuse
- ◆ A system for tracking both the data and the media while stored at the off-site location
- ◆ A method for reporting on media shipments
- ◆ Efficient resource use (where resources comprise tapes, drives, media servers, and network)

How Vault Meets these Requirements

- ◆ By using standard NetBackup utilities, thus becoming a natural extension of NetBackup, capitalizing on the backup duplication, backup tracking, and media tracking facilities which are already familiar to the user;
- ◆ By storing the required information within NetBackup and Media Manager catalogs;
- ◆ By using profiles to specify configuration parameters such as which backups to send to the vault, whether or not to duplicate the backups, which media servers to use, which volume pools and volume groups to use;
- ◆ By vaulting original or duplicated backups;
- ◆ By backing up the NetBackup and Media Manager catalogs and sending that backup to the vault with the rest of the media;
- ◆ By supporting media eject and inject for Media Manager robotics;
- ◆ By providing the option to duplicate backups locally or across the network, so as to offload important backup resources.

What Vault Does

NetBackup Vault extends the NetBackup infrastructure to provide automated duplication of specific backup images, and automates the process of off-site media rotation (a critical component of any backup or disaster recovery strategy). NetBackup Vault manages off-site storage and retrieval of media for original backups, duplicate backups, and catalog backups. Additionally, NetBackup Vault generates reports to track the location and content of each piece of media. Vault uses existing NetBackup functions for all operations, such as duplication of backups, media control, reporting, and the ejecting and injecting of tapes used for vaulting. Information from Vault is integrated with other NetBackup components and appears in the NetBackup Activity Monitor.

Vault is a three step process:

1. Duplication of backups
2. Backup of NetBackup catalog
3. Ejecting of media/reporting

For a given group of backups, you can perform any or all of the above steps.

How Vault Fits into the NetBackup Product

NetBackup Vault provides extended functionality to the core NetBackup product. It relies on NetBackup services and the NetBackup catalog for its information, and maintains a backup of the image catalog for recovery services. For example, Vault:

- ◆ Uses existing Media Manager services for fundamental media and robotic management
- ◆ Uses the NetBackup catalog and the Media Manager database to keep track of which images have already been vaulted
- ◆ Uses the Media Manager database to request that expired media be rotated back into the robot for reuse
- ◆ Uses Activity Monitor to display status of the vault and duplication jobs involved
- ◆ Stores information pertaining to vault in dedicated fields to the NetBackup volume database
- ◆ Provides access to the Inline Tape Copy (multiple copies) feature
- ◆ Works with the NetBackup logs, and provides its own logs under NetBackup/Logs
- ◆ Provides email notification of the status of Vault jobs
- ◆ Uses the same type of role-based security that is used by NetBackup
- ◆ Provides multi-eject support
- ◆ Configures Vault policies through the policy management GUI.
- ◆ Is subject to some of the NetBackup scheduler's prioritization rules, although Vault is not yet fully integrated into the scheduler.

How to Access Vault

NetBackup Vault is installed on a NetBackup 4.5 master server. The product consists of:

- ◆ Java GUI (for UNIX systems) which appears as a node on the NetBackup Console



- ◆ Windows GUI (for NT4 and W2K) which appears as a node on the NetBackup Console
- ◆ Executables that reside on the NetBackup server and are turned on with a license key
- ◆ Menu-based user interfaces
- ◆ Command line utilities

Note *For legacy NetBackup Vault users only:* Vault no longer operates through editable scripts and parameter files. To configure Vault, use either the graphical user interface (GUI), or the menu-based user interface (vltadm). Parameter files have been replaced with Profiles, which you will configure through the GUI or the MUI. For more information about profiles, please see “Configuring a Profile” on page 51.

What is Vaulting?

Vaulting is the process of choosing backup images to duplicate or eject, optionally duplicating backups, ejecting duplicate or original media, storing it at an off-site location, and later returning expired media to your robot. We use the term ‘vault’ to refer both to a logical entity associated with a particular robot, and to the physical storage location of a set of tapes off site. A vault session is the process of vaulting a particular collection of data.

Vaulting Duplicate or Original Media

Vault distinguishes between ‘original’ media and duplicate media as follows:

- ◆ Original media is any media for which the images were created by backup policies. This includes all copies created by a multiple-copy configuration in a backup policy.
- ◆ Duplicates are copies created after the backup policy completes. Typically this is done by the duplication of a Vault profile.

Vault lets you store either the original or the duplicate off site. Sometimes you may want to send the original backups off site instead of creating duplicates to send off site. For example, if limited resources mean that the time allotted for performing a duplication is shorter than the time the task actually requires, you may decide to send some of the original backups off site.

To specify which backups to vault, you create a profile which contains the criteria used to select the backups.

What Is the Primary Copy?

The primary copy is the copy designated for use by the restore process. This process always uses the primary copy, so you will want to ensure that the copy that will remain in your robot is primary.

NetBackup's backup process can make up to ten copies. By default, the primary copy is the first (or only) copy made by the original backup process. If you do not configure the Vault profile to change the primary copy, then the copy that was designated primary before the Vault session runs remains the primary copy. If you wanted to eject and vault that first copy, you could configure the Vault profile to specify a different primary copy.

What is Duplication?

Duplication is the process of creating another copy of a backup. With NetBackup Vault 4.5, you can make up to four copies of a backup at a time and designate any of those copies as the primary copy, or you can skip the duplication step and configure Vault to vault an original backup image.

The primary copy is the copy designated for use by the restore process. Duplication always uses the primary copy, so you will want to ensure that the copy that will remain in your robot is the primary copy.

A special type of duplication, Inline Tape Copy, is the process of creating more than one tape copy of a backup simultaneously. You can do this at either of two stages:

- ◆ If you know, when originally creating a backup, that you will want more than one copy, you can create multiple copies through NetBackup Management.

At this stage, you specify the number of multiple copies in the policy or schedule configuration.

- ◆ If you decide, after having created your backups, that you want more than one copy, you can create them through Vault Management.

At this later stage, you specify:

- The number of copies
- The volume pool for each copy
- The destination storage units

from the Catalog node, on the NetBackup Console, the command line, or a Vault profile.

At this stage, you specify that one of the new copies will become the primary image (primary copy). You may want to do so when the first copy made by the original backup process is the copy to be ejected and vaulted, and a duplicate copy is to remain on site.



What is a Catalog Backup?

The NetBackup and Media Manager catalog consists of internal databases of information about the NetBackup configuration and any backups that have been performed. The information about backups includes records of the files and the media on which the files were stored. The catalogs also have information about the media and storage devices that are under the control of Media Manager.

Creating a catalog backup is a different procedure from creating a backup of other data. You must vault a catalog backup with your backup data to efficiently recover from a disaster.

What is the Eject Step?

Eject is the process of physically ejecting media from the robot. Media that are ejected are tracked by Vault reporting facilities and will be recalled from the Vault vendor for reuse after the media expire. Media can be ejected automatically by the scheduled Vault job or manually after the job has completed. Media can be ejected for each job individually or can be consolidated into a single eject operation for multiple Vault jobs.

What is the Reporting Step?

Reports are the mechanism the user and the vault vendor have for tracking vaulted media. There are four types of reports:

- ◆ Media being sent to vendor
- ◆ Media to be returned from vendor
- ◆ Various inventory reports
- ◆ Recovery report, which identifies media that must be returned from the Vault for efficient disaster recovery

Reports can be generated automatically by the scheduled Vault job or manually after the job has completed. Reports can be generated for each job individually or can be consolidated with a consolidated eject operation.

Terminology

Alternate Read Server

The *Alternate Read Server* is the server used to read a backup image which was originally written by a different media server. The media server specified as *Alternate Read Server* must have access to the media containing the backup image or images it is configured to read.

Consolidated Eject

A *Consolidated Eject* is the process of ejecting media for more than one Vault session at a time. A *Consolidated Eject* can be performed for one or more logical vaults at one time.

Consolidated Report

A *Consolidated Report* is the process of generating reports for more than one Vault session at a time. A *Consolidated Report* can be performed for one or more logical vaults at one time. Consolidated reports are organized by report title, not by vault.

Destination Storage Unit

The *Destination Storage Unit* is a storage unit to which Vault sends the data from a duplication operation. If the duplicated backup images are to be vaulted, then the destination storage unit must correspond to the robotic volume group.

Off-site Volume Group

The *Off-site Volume Group* is the volume group in which media will appear after having been ejected from the robot for vaulting. When Vault ejects media it is moved from the robotic volume group to the off-site volume group.

Off-site Volume Pool

An *Off-site Volume Pool* is a volume pool containing media that is to be ejected and vaulted. Backup images written to an off-site volume pool by an original NetBackup backup policy or by Vault's duplication feature will be ejected and vaulted. More than one off-site volume pool can be specified for the Eject step of a Vault profile.

Original Backup

An *Original Backup* is a backup image created by a backup job. A single backup image or all backup images created by an Inline Tape Copy (multiple copy) configuration are considered original backups. A backup image created by a duplication job is not an original backup.

Robotic Volume Group

The *Robotic Volume Group* is the volume group from which media will be ejected and vaulted. When Vault duplicates backups, they are duplicated to media in the robotic volume group.



Source Volume Group

The *Source Volume Group* is the volume group from which Vault can select backups to duplicate. This parameter is used to restrict the list of backups from all backups that reside on media in any volume group to backups that reside on media in a single volume group. Where a volume group corresponds to a particular robot, the profile will duplicate only backups on media in that robot. The *Source Volume Group* is normally only specified if you have multiple robots attached to the same server, for example you want to duplicate backups that reside in robot 0 to media that reside in robot 1.

Unassigned Media

Unassigned Media are media that contain no valid images. A piece of unassigned media has an entry in the volumes database but no entries in the images database. *Unassigned Media* do not have a time assigned in the Media section of the GUI.

Installing NetBackup Vault 4.5

2

This chapter outlines the steps required to install NetBackup Vault 4.5 on both UNIX and Windows systems. The following topics are covered:

- ◆ Supported Platforms
- ◆ Supported Robots
- ◆ Installation Prerequisites
- ◆ Installing
 - Installing on UNIX Systems
 - Installing on Windows Systems
- ◆ Upgrading
- ◆ Uninstalling
 - On UNIX Systems
 - On Windows Systems

Supported Platforms

NetBackup Vault supports the same operating systems and versions as NetBackup with the exception of the NCR and Sequent operating systems, which are not supported by Vault 4.5. Additionally, the Vault Java GUI does not support IRIX.

For more information about supported platforms, please see the NetBackup *Release Notes*.

Supported Robots

Vault supports Storage Tek's ACSLS, and Media Manager's TLD, TL8, TLH, and TLM architectures for robotic control. Robots must have a media access port (MAP).



Installation Prerequisites

- ◆ NetBackup must be up and running on the server.
- ◆ Vault must be installed on a NetBackup master server.
- ◆ You must have a valid license key for Vault.

Installing on UNIX Systems

▼ To install Vault on a UNIX server

1. Log in as the root user on the master server.

If you are already logged in, but not as the root user, execute the following command:

```
su - root
```

2. Verify that the Vault license is present, and add it if it is not. Enter the following command to list or add keys:

```
/usr/netbackup/bin/admincmd/get_license_key
```

3. Insert the CD-ROM containing the Vault software into the drive.

4. Change the working directory to the CD-ROM directory:

```
cd /cd_rom_directory
```

where *cd_rom_directory* is the path to the directory where you can access the CD-ROM. On some platforms, you may need to mount this directory.

5. To install NetBackup Vault, execute the following:

```
./install
```

- a. Select the **Add-On Products and Database Agents** option.

As other NetBackup products are included on the CD-ROM, a menu appears.

- b. Select the NetBackup Vault option.

- c. Enter **q** to quit the menu.

- d. When asked if the list is correct, answer **y**.

Vault is installed in `/usr/opensv/netbackup/db` and `/usr/opensv/netbackup/bin`.



6. Once you have installed Vault, you should configure the email address for notification of sessions status and enter alternate media server names if appropriate. See “Entering Addresses for Email Notification” beginning on page 92 and “Adding Alternate Media Server Names” beginning on page 95.

Installing on Windows Systems

Vault is installed as part of the NetBackup installation, but you cannot use it until it is licensed. If Vault was included as part of your base NetBackup license, you will need no additional license keys. However, if Vault is not included in the base NetBackup license, you will be prompted for a license key. Once you supply the key, Vault will be available for use.

▼ To add the Vault license key

1. From the NetBackup Administration console, highlight the **Help** menu.
2. Choose **License Keys**.
3. From the Current Licenses screen, click the **New** button.
The Add a New License Key dialog displays.
4. Type the Vault license key and click the **Add** button.
The license key appears on the Current Licenses screen.
5. Once you have installed Vault, you should configure the email address for notification of sessions status and enter alternate media server names if appropriate. See “Entering Addresses for Email Notification” beginning on page 92 and “Adding Alternate Media Server Names” beginning on page 95.

To complete the installation, you must configure appropriate NetBackup attributes for use by Vault and identify which NetBackup policies you wish to vault. For configuration information, please see “Configuring Vault” on page 45.

Upgrading

This is the first version of Vault which begins integration into the NetBackup product. Legacy users should speak with a VERITAS consultant about upgrading to the integrated NetBackup Vault 4.5 release. If you have a simple bpvault configuration, see “Upgrading bpvault to NetBackup Vault 4.5” for guidance.



Uninstalling

If you choose to permanently remove Vault from your NetBackup installation, identify other configuration items specific to Vault, and remove them as well if they are no longer in use, for example, Vault volume pools or policies.

On UNIX Systems

To remove Vault from your NetBackup installation, follow these steps.

▼ To remove Vault

On the master server where you initially loaded the Vault software, execute the following:

- For Solaris NetBackup servers, execute:

```
pkgrm VRTSnbvlt
```

- For other UNIX NetBackup servers, delete the following:

```
/usr/opensv/netbackup/bin/bpbrmvlt
```

```
/usr/opensv/netbackup/bin/vlt*
```

```
/usr/opensv/netbackup/bin/goodies/vlt*
```

```
/usr/opensv/netbackup/vault
```

To remove the Vault database, also delete:

```
/usr/opensv/netbackup/db/vault
```

To delete Vault logs, delete:

```
/usr/opensv/netbackup/logs/vault
```

On Windows Systems

To remove Vault from your NetBackup installation, simply remove the license key.

▼ To remove the Vault license key

1. From the NetBackup Administration console, highlight the **Help** menu.
2. Choose **License Keys**.
3. From the Current Licenses screen, highlight the Vault license key.

Caution If Vault was included as part of the base product key, performing Step 4 will delete your base key, and you will be unable to use NetBackup. If you do not want to delete the NetBackup license key, do not continue.

4. If you want to continue with the deletion, click the **Delete** button.

The license key disappears from the Current Licenses screen, and Vault Management disappears from the console tree.





This chapter describes the types of data to collect from NetBackup before you start using Vault. It also contains guidelines to help you in planning your Vault configuration. Reading the recommendations below will help you to make efficient use of your NetBackup resources.

This chapter covers the following topics:

- ◆ Notes and Recommendations
- ◆ Managing Multiple Media Types within a Vault
- ◆ Collecting NetBackup Information
- ◆ Vault Policy Types
- ◆ Gathering NetBackup Information
- ◆ ACSLS Robot Information
- ◆ Media Manager Configuration
- ◆ Network Configuration for Receiving Reports
- ◆ Vault Configuration

Note This chapter assumes you are familiar with basic NetBackup concepts, such as policies and storage units. If not, please refer to the *VERITAS NetBackup System Administrator's Guide*.

Notes and Recommendations

Managing Multiple Media Types within a Vault

A vault is a logical entity which refers to a collection of tape drives within a robot. A profile is a set of configuration instructions, chosen by the storage administrator, and carried out in a vault session. A profile is associated with a particular vault. You can manage multiple media types within the same vault by creating more than one profile.



For example, if some number of policies should be written to HCART for off-site storage, while a different set of policies should be written to DLT cartridge, and the original backup images were all written to a single media type (for example, DLT cartridge), you would create a profile for each set of policies.

First, define the set of policies to duplicate. Then define the storage unit that should be used for the destination copy of the image. If both types of drives are attached to the same NetBackup server, then these fields would be the only differences between the two profiles. If all of the images within both sets of policies were backed up on the same NetBackup server, the source server field would be the same.

This situation would also apply to multiple media densities (e.g. DLT4000 vs. DLT7000, 4890 format vs. Timberline format, etc.), since NetBackup would recognize the media differently (for example, DLT vs. DLT2, HCART vs. HCART2).

If the duplicate images must be stored on the same media type as the original images, and the NetBackup media server possesses enough resources (that is, at least two drives of the media type), then you do not have to create separate profiles. For example, if the media server is connected to two DLT tape drives and two Timberline drives, and some set of backup images has been written on both DLT and half-inch cartridges, NetBackup can handle the mount requests and assignment of tape drives automatically within the bpduplicate interface.

Collecting NetBackup Information

You must identify NetBackup volume pools, volume groups, servers, and policies for use with Vault.

Note For procedures on how to create these objects in NetBackup, please refer to the *VERITAS NetBackup System Administrator's Guide*.

▼ To prepare NetBackup for use with Vault

1. You must designate several volume pools specifically for use with Vault. You can create new volume pools, or you can use volume pools that already exist in NetBackup as long as the pools are for Vault's use exclusively.
 - Create or specify a volume pool to use for duplicating backup images. You can choose any name for the pool, for example, Off-site Backups.

It is best to avoid using the NetBackup volume pool for data you want to send off site. This is because the NetBackup volume pool is the default volume pool, which means it's very likely that you would end up sending far more data off site than necessary.

- Create or specify a volume pool for Vault's NetBackup catalog backups. This pool is only for catalog backups. Choose an easily identified name for this pool, for example, Catalog_Off-site.

Note The catalog backup volume pool must be a different pool than those used for duplicating backup images because NetBackup will freeze the duplication media if it detects a catalog header on the tape.

Note NetBackup does not automatically move media to the catalog backup pool from the scratch pool. You must monitor the media in this pool for use. There must be a tape in the catalog volume pool before the Vault session runs. If there are no available tapes in the catalog volume pool, the session will fail.

- Identify the volume pools that NetBackup uses for backups.
 - Identify the policies and schedules NetBackup uses.
2. When you create a vault, you must select the volume group that the media in the robot belong to. The media in this volume group are the media Vault will use if you choose to duplicate backups or the NetBackup catalog.
 3. Choose a simple policy to use as a Vault test run.

Once you have configured Vault, you can use this policy to make sure you get the results you expect from your configuration in a preview of results from each Vault profile.

Vault Policy Types

Determine the types of off-site vaulting policies you will require to meet your backup needs. For instance, are you able to duplicate all of your data within a specific window, or do you need to send original media off site? Do you wish to duplicate the backups for a group of critical servers, and send original media off site for the other servers? Do you wish to make multiple copies at one time? Do you want to back up images to disk and make the copies at some other time?

You must designate at least one Vault policy.



Off-site Policy Considerations

For currently defined off-site policies, as well as for any new off-site policies you may create, list the following information. A particular off-site policy may be designated for more than one vault. List the following information for each vault created.

Name of off-site vault vendor	List the name of your off-site vault vendor (for example, Arcus or Iron Mountain).
Starting Slot ID	Identify the number of the first off-site slot number (seed number) for each vault.
Schedule to pick up/return tapes	The schedule agreed upon between you and the off-site vendor determines how often you will run an eject/report session. Some sites send media off site daily, others weekly.
Frequency of off-site vaulting	Consult your off-site schedule to determine how often you need to run an eject/report session.

Gathering NetBackup Information

Use the tables below to help you organize the information you collect from NetBackup. Refer to your current NetBackup configuration to gather the information.

Master Server, Media Servers, and Robotic Devices

For the master server, list the number of media servers. For each media server, list attached robotic devices and storage units.

Server Host Name	The name of the master server where Vault is installed.
Operating System Level of Master Server	The release of the operating system on the master server, such as Windows 2000, service pack 2.
Number of Media Servers	The number of media servers associated with this master server.
Operating System Level of Media Servers	The release of the operating system on the media server or media servers.

Types of Robotic Devices	The robotic devices associated with the media server. Use the appropriate NetBackup terminology (for example, TLD, ACS, TL8) or specify the actual hardware manufacturer and model names for each device
Storage Unit Name	Specify the NetBackup storage units that are associated with each media server. You can use the command <code>bpstulist -U</code> to generate a printout of existing storage units.
Number of Drives	The number of drives in each storage unit.
Robot Number	The robot number this storage unit is associated with, as defined in Media Manager.
Robot Control Host	The name of the media server that controls the robot.

NetBackup Policies

Collect information for each schedule/policy pair you may want to consider for off-site rotation:

NetBackup Policies

Policy Names	All policies that you want to consider for off-site rotation. For policies that affect critical servers that you wish to duplicate, list "Duplicate". For policies that cover too much data to duplicate, specify as "Original". You can use the command <code>bppllist</code> to get a full list of policies.
Schedule Names	If multiple schedules exist within a policy, and you want the schedules to be part of the off-site policy, list each policy/schedule pair separately. For example, you may want to send only duplicates of full backups off site, but send the original incremental backups off site.
Off-site Original, Duplicate, or Both	Which type of backup is being vaulted for each policy/schedule combination
Storage Unit	The storage units for each policy.
Retention Period	The retention period for each schedule so that you will have an idea of when to expect the off-site media to return.



Robot Information

Collect the information below for each robot.

Robot Number	The robot number assigned by Media Manager.
ACSLS Server	The name of the ACSLS server.
ACS Number	The corresponding ACS number for this robot. Vault also requires ACSLS configuration information. You can obtain this information through Media Manager's <code>tpconfig</code> or directly from the ACSLS console by using ACSLS commands such as <code>query acs all</code> or <code>query lsm all</code>
LSM Number	The corresponding LSM number for this robot. Vault also requires ACSLS configuration information. You can obtain this information through Media Manager's <code>tpconfig</code> or directly from the ACSLS console by using ACSLS commands such as <code>query acs all</code> or <code>query lsm all</code>
CAP Capacity	The capacity of the Cartridge Access Port. You can obtain this information by using the ACSLS command <code>query cap all</code> from the ACSLS console.
CAP Numbers	The identifiers for the Cartridge Access Port. You can obtain this information by using the ACSLS command <code>query cap all</code> from the ACSLS console.

Media Manager Configuration

Identify the volume pool names and their respective media density/type (for example, DLT, or HCART) for each type of vault job: catalog backups, duplicates, or originals. Specify Duplicates for media that will be allocated for duplicated images and Off-site_Full for media that contain the actual backup images. Vault also allows you to define a pool of tapes that will be used for backing up the NetBackup catalogs and then sending them off site. If you wish to use this feature, define a separate pool of tapes that will only store these off-site copies of the NetBackup catalogs, and list them as Off-site NBU Catalog Backups.

Duplicates	Originals
Volume pool or pools; media types	Volume pool or pools; media types

Network Configuration for Receiving Reports

Vault generates reports each time a vault job is processed. To automatically print these reports, or to have them delivered in an email, you will need the following information:



Note All reports will be sent from the NetBackup master server, so the mail addresses and network printers listed should be configured on the master server prior to Vault configuration.

Type of reports desired:	Refer to the Reporting chapter to select the report types appropriate to your installation.
Report destination	You can choose to send the report to email addresses, to save it to disk, and/or to print it.
Report mode	Indicate if you want Vault to automatically produce the reports, or if you want to run them manually through <code>vltopmenu</code> . Some reports require that the media be ejected before Vault can generate them.

Vault Configuration

A vault is a logical entity which refers to a collection of tape drives within a robot. A profile is a set of configuration instructions, chosen by the storage administrator, and carried out in a vault session. A profile is associated with a particular vault.

Identify the information below for each vault. You must define at least one logical vault for each robot. For example, if your NetBackup configuration contains three distinct TLD robots (not connected with pass-through devices), you would define at least three logical vaults, one for each TLD robot.

Vault Identifier	Assign a name or number to this vault.
Originals or Duplicates	Whether this vault is for original backups, or duplicates.
Off-site Slot Seed	The slot ID of the first slot in the off-site vault.
Off-site Vendor	The off-site vault vendor.
Off-site Volume Group	The name of the off-site volume group.
Robotic Volume Group	The volume group or groups associated with the robot.

Profile Configuration

Identify the purpose of each profile you configure, and the resources for a vault session using this profile.



The first section of profile configuration corresponds to the first tab of the Vault Configuration dialog box.

Profile Name	A name for the profile which reflects the purpose of the profile.
Type of Backups	The types of backups (full, incremental, etc.) this profile will capture. This is an optional criterion.
Time Period for Selected Backups	The period of time from which the profile will select backups.
Clients to Consider for Vaulting	The clients you want to select backups for.
Media Servers to Consider for Vaulting	The media servers you want to select backups from.
Policies to Consider	A list of policies that you want to use to select backups. These policies and schedules are based on the storage unit used for backups. Since the storage unit is related to a specific robot number, vault the policies and schedules by the robotic device.
Schedules to Consider	A list of schedules you want to use to select backups. Note These policies and schedules are based on the storage unit used for backups. Since the storage unit is related to a specific robot number, group the policies and schedules by robotic device.

The next collection of information is for use with Duplication sessions. You will use this information on the Duplicate tab of the Profile Configuration dialog box.

How many drives per media server do you want to use for Vault sessions? You can specify how many drives to use for each storage unit.

Storage Unit Name	The storage unit associated with each media host. Consider how many drives in each storage unit you want to use for vault sessions. You may choose to keep some drives available for restores or backups while duplication is running.
--------------------------	---

Server Host Name	<p>The name of each media server that controls the drives you want to use for the vault process. This server should also be bound to a storage unit within the NetBackup configuration. Most sites will define all drives (of a given media type) that are attached to a server as one storage unit, since this is recommended during the installation of NetBackup.</p> <p>For example, two servers may share drives within a robot. Each server is connected to four drives, which means the robot houses eight drives. If you only wish to use two drives on each server for duplication (so that other operations such as restores can utilize the remaining drives), then the total number of drives that would be used for duplication is four, or two pairs.</p>
-------------------------	---

The collection of information below is for use with the Report/Eject/Catalog Backup tabs:

Reports Generated	The Vault reports you want this profile to generate.
Volume Pools to Use	<p>For vaulting duplicates, the destination volume pool or pools. For a single copy, specify one pool. For multiple copies, specify multiple pools. Each unique copy should have a unique pool. (Duplicate tab)</p> <p>For vaulting originals, this is the eject pool or pools. (Eject tab)</p>
Suspend Option	For sessions which vault originals, do you want to configure the suspend option?
Automatic or Manual Eject	You can configure Vault to automatically eject all media upon completion of the vault session, or to allow the operator to process the media ejects. For example, if a robot has only one export slot, and there are multiple tapes to be ejected, the eject process must be handled manually. For this option, select automatic or manual.



Specify the information below for catalog backups.

Backup NetBackup Catalog?	<p>If so, specify the server and volume pool to use for the catalog backup. If you want to send a copy of the NetBackup catalogs off site during a vault session, you must decide which vault will create this off-site copy. Only one vault should be responsible for creating an off-site copy of the NetBackup catalogs. Normally, the robot that is attached to the master server is the choice for creating a backup of the NetBackup catalogs, since the master server is usually used for creating an on-site copy of the NetBackup catalogs. (See the discussion of NetBackup catalog backups in the <i>NetBackup System Administration Guide</i>).</p> <p>The host name of the server responsible for these backups should be listed, along with the volume pool that is reserved for off-site NetBackup catalog backups. Multiple backups of the NetBackup catalog backups can be created, but most sites make one off-site copy. Since NetBackup does not track the NetBackup catalog backups within the catalog, you must specify how long to retain these backups off site. Few sites retain these backups longer than 10 or 20 days, as the data becomes obsolete quickly.</p>
Number of off-site copies	How many copies of the catalog backup do you plan to send off-site?
Retention period for catalog backups	How long do you want to keep a valid copy of the catalog backup in the vault?
Multiple-tape catalog backups	If your catalog requires the NetBackup two-step process, you must specify the backup policy to use.

This chapter provides best practices and suggestions relating to the following broad areas:

- ◆ Preferred Strategies
- ◆ Make Sure You Vault All You Want To
- ◆ Do Not Vault More Than You Want To
- ◆ Ensuring Report Integrity
- ◆ Preparing for Efficient Recovery
- ◆ Avoid Resource Contention During Duplication
- ◆ Avoid Sending Duplicates Over the Network
- ◆ Increase Duplication Throughput



Preferred Strategies

Preferred Strategy: Vault Original Backups

Every Vault profile can vault original backups, duplicated backups, or both. We strongly recommend vaulting original backups, for several reasons:

- ◆ Vaulting originals uses fewer drives than duplicating backup images from the original tapes. For instance, if your backup job creates two copies of the backup, two drives are needed. On the other hand, if your backup job creates one copy of the backup and your vault job subsequently creates one duplicate of the backup, then the original backup job consumes one drive, and the subsequent vault job consumes two drives (one read drive and one write drive). Over time, duplicating backup images consumes more drive-time than writing multiple copies of the backup during your original backup job.
- ◆ Vaulting originals avoids drive contention and media contention problems with concurrent duplication and backup processes. If you have a great deal of time during which no backups are running -- time during which Vault can do duplication -- such resource contention may not be an issue. But as your data needs grow, you may not continue to have this luxury.
- ◆ Vaulting originals avoids the complexity of configuring for duplication. If your backup environment uses multiple media servers and/or multiple robots, or if you require different retention periods for the vaulted copies of different types of data, it can be difficult to configure the duplication step of your Vault profiles. If you are not running in a SAN environment, it is easy to unwittingly send large amounts of data over the network. It is easiest to just avoid the duplication step altogether.

For these reasons, we strongly advise you to make multiple copies of your backups during the original backup jobs, and vault one or more of these original backup copies. If you do so it is important that you read and fully understand the sections on using the suspend option and Vaulting Original Backups in a 24x7 environment.

Disk Staging

Another excellent way to avoid the complications that arise when trying to share resources (drives and media) with backup jobs is to use disk staging for your backups. Because disk is fast and less expensive than tape drives, it is often advantageous to send your backups to disk. Later, when your backups have completed, schedule your Vault sessions to duplicate the original disk backup to two (or more) tapes, one on-site tape, and one off-site tape. The Vault profile can be configured to automatically free up the disk space for the next round of backups. A disk image will not be expired if the duplication operation did not succeed for that image.

Advantages:

- ◆ Backup window is shortened.
- ◆ Minimizes tape drive usage. Sending the original copy to tape, then duplicating to a second tape requires one drive to make the first copy and two drives (a read drive and a write drive) to make the second copy.
- ◆ Eliminates drive contention and media contention between backup jobs and vault jobs which do duplication.

Make Sure You Vault All You Want To

Overlap the Time Window in the Profile

The Vault profile uses a time range as a part of the criteria for choosing the backup images to be vaulted. Configure this time range to be at least one week longer than the frequency at which this profile is to be executed. For instance, suppose you have a profile which does duplication and which you will run daily. In this case the time window should extend from 7 days ago to 0 days ago.

Suppose you have a catastrophic hard failure on one of your media servers which takes 4 days to repair. The next time your Vault profile runs, it will automatically vault backups created before the 4-day downtime. If you only execute the profile once a week, make sure the time window encompasses at least two weeks. That way, if hardware problems cause the process to fail one week, the backups will be picked up the next week.

Vault will neither duplicate nor eject a backup image that already has a copy in the off-site volume group. Therefore, when vault sessions are successfully completed on a daily basis, only backups from the previous 24 hours will be processed.

Adding 7 days to the time window will cause Vault to process a larger list of images. This will consume a bit more processing time, but due to the batch nature of Vault, this extra processing time is not a problem in many environments.

Consequences of Not Overlapping the Time Window: Missing Data

When a vault session gets delayed, some backup images may be missed if the time window does not allow Vault to pick up images from a wider time range. For instance, suppose the time window for your daily profile extends from 1 day ago to 0 days ago. On Tuesday, the robot has mechanical problems and the vault profile is unable to run. Consequently, Monday night's backups are not vaulted. On Wednesday, the robot is fixed. When the next vault session begins on Wednesday, it will only pick up backup images that



were created during the previous 24 hours; so Monday night's backups are still not vaulted. If the profile's time window had spanned more than 1 day, the session would have picked up both Monday night's and Tuesday night's backups.

Resolve Multiple Names for a Single Server

For every media server, you should add an entry on the Alternate Media Server Names tab of the Properties dialog. At a minimum, there should be, for each media server, an entry which contains both the abbreviated name and the fully qualified name. Also add any other names by which a media server has ever been known. Taking this action will avoid a number of problems. For example, if you do not list alternate names for media servers, some images may not be recognized to match the Choose Backups criteria and may therefore not get vaulted.

If you have multiple NIC cards in your server, make sure that the server name or IP address associated with each NIC card is listed in the Alternate Media Server Names dialog.

Specify Robotic Volume Group When Configuring Logical Vault

If you want a tape to be ejected, make sure it is in the robotic volume group and in one of the off-site volume pools specified for the Eject step.

The Robotic Volume Group is specified when you configure a logical vault. This is the volume group in the robot. Consequently, this is the volume group from which media will be ejected. When the tapes are ejected, Vault moves the media IDs to the off-site volume group specified for that vault.

A tape will only be ejected if it is in the robotic volume group *and* in one of the off-site volume pools specified for the Eject step.

Multiple Volume Groups (Multiple Robots)

A Vault profile will only eject tapes from the robotic volume group specified for the vault to which the profile belongs.

Volume groups cannot span robots. Typically a volume group identifies a specific robot.

The backup images selected by the Choose Backups criteria can come from multiple volume groups (multiple robots). This is useful if your profile is to do duplication, copying backup images from multiple robots to the robot which will do the ejects. This must be configured with care as described in "Alternative A: Single Dedicated Robot for Vault Processing" on page 35.

If your profile does not do duplication, then it is most efficient to set the Source Volume Pool on the Choose Backups tab to be the same as the Robotic Volume Group configured for the vault.

If no tapes are ejected, it might be because:

- ◆ All images have already been vaulted. Vault determines that a backup image has already been vaulted if it already has a copy that is on a tape which is in the off-site volume group.
- ◆ The tapes to be vaulted are in a volume group other than the robotic volume group specified for the vault to which the profile belongs.

Ensuring Report Integrity

A best practice is to decide up front whether you want your Vault reports to group media by robot or by vault. Based upon this decision, you will relegate a different off-site volume group name for each robot or for each vault.

With the exception of reports for media going off site which pertain to a particular profile, Vault reports are generated based on the off-site volume group. Therefore, if vaults for multiple robots share a single off-site volume group, reports that are intended for only one robot could also include media that correspond to other robots.

Integrity of Media Organized by Robot: To guarantee that no media from one *robot* will appear on the report for another robot, specify a unique off-site volume group for the vaults within each robot.

If your chief concern is to maximize the reuse of tapes, use the same off-site volume group for all logical vaults within a robot. Reports will not seem consistent for an individual logical vault, but this strategy will maximize the frequency with which tapes are returned for reuse. Every time the *Picking List for Vault* report is generated for any profile within any vault for the robot, tapes from all profiles and logical vaults for that robot could be recalled for reuse (depending on how profiles share off-site volume pools; see discussion in next section).

Integrity of Media Organized by Vault: Likewise, if you decide that you would like to guarantee that no media from one *vault* will appear on the report for another vault, specify a unique off-site volume group for each vault.

If your chief concern is consistency, use a different off-site volume group for each logical vault within a robot. A single report will contain media from only a single logical vault.

For instance, the *Picking List for Vault* report identifies media to be returned to you for reuse. If you have more than one robot, the integrity of this report ensures that the correct media get returned to the correct robot. This report is based on the off-site volume group



of the vault to which the profile belongs. If some other logical vault from a different robot used the same off-site volume group, then the *Picking List for Vault* might include media that should actually be have been returned to a different robot.

Consequences of Sharing an Off-site Volume Group Across Multiple Robots

If profiles from multiple robots share both an off-site volume group and one or more off-site volume pools, then your vault vendor will return a group of tapes (for a single *Picking List for Vault* report) that were ejected from multiple robots. The operator will need to identify which tapes should be injected into each of the robots. If mistakes are made through this manual process, you will end up with the wrong media (and possibly the wrong number of media) in your robots.

Further Delineating Reports

As discussed above, Vault uses the off-site volume group to query for the media to include in the reports. It also uses the off-site volume pools listed for the eject step of the profile for the same purpose. Therefore you can use either the off-site volume group or the off-site volume pool(s) to select media for each robot, vault, or profile.

Separating Reports by Profile: If you want the reports to include only media for a single profile, you must specify a unique off-site volume pool for every profile.

Separating Reports by Vault: If you want to separate reports for each vault, you could specify a common off-site volume group for all vaults within a robot, and specify a unique off-site volume pool for all profiles within each vault. Doing so ensures that each report would contain media from a single vault. In addition it minimizes the number of off-site volume groups in your configuration

Making a Report Match Media Assigned to a Logical Vault

If you want the Inventory List for Vault report to exactly match the media assigned to a particular logical vault (stored for one group of slots), then you can choose either of the following strategies:

- ◆ Use off-site volume groups to delineate reports: Use a unique off-site volume group for each logical vault, and share off-site volume pools across all profiles, all vaults, and all robots as desired.
- ◆ Use off-site volume pools to delineate reports: Use unique off-site volume pools for profiles within each logical vault (do not share off-site volume pools across logical vaults), and share the off-site volume group across vaults and reports as desired.

Do Not Vault More Than You Want To

Send Only the Intended Backups Off-site

When configuring your backup policies, never send a backup that you do not intend to vault to a volume pool which is to be used for vaulting.

If a number of images reside on a given piece of media, and these images are created by multiple backup policies, Vault will send this media off site even if only one of the policies is specified in the profile. For example, assume the tape ABC123 has six images on it, and that those six images were backed up from policy 1 and policy2. Assume that you have only specified policy1 in the profile. The tape with the media ID ABC123 would be sent off site, even though it also contains images from policy 2.

Vault is largely controlled by volume pool. Ultimately, no backup image will be vaulted if it is not on a tape that is in the appropriate volume pool.

Choosing Which Images for Which Media

When working with originals, you should not put images that you want to keep on site onto the same media as images that you want to vault, because you will end up vaulting all of the images on that media.

If a number of images reside on a given piece of media, and these images were selected by multiple NetBackup policies, Vault will still send this media off site even if only one of the policies is specified in the profile. For example, assume the tape ABC123 has six images on it, and that those six images were backed up from policy1 and policy2. Assume that you have only specified policy1 in the profile. The tape with the media ID ABC123 would still be sent off site, even though it also contains images from policy2

Sending Backup Images to Volume Pools

Use different volume pools for backup images you want to keep on site and backup images you want to send to the vault. If you use the same volume pool for both, you will vault the backup images that you meant to keep on site.

In addition, if you use the same volume pool for both, a deadlock situation may result if your vault profile is doing the duplication step, because the NetBackup duplication process will attempt to read the backup image from the same tape to which it wants to write the image.

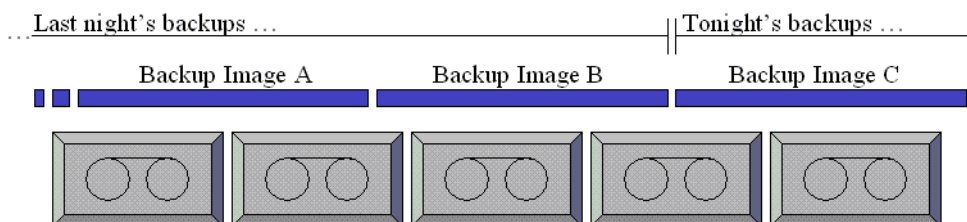


Use the Suspend Option to Avoid Vaulting Partial Backups

When a tape is suspended, no more backups will be written to the tape until the images on it expire. Suspend is used when vaulting original backups (i.e. when Vault is not doing duplication). Suspension write-protects original backups on the specified media so that when the media are vaulted, they will not contain fragments of newer images.

Since images are usually ejected in batches, there may be backups that you would like to vault but cannot because they are not complete by the time you want to start vaulting. Therefore, most tapes begin with a partial backup image and end with a partial backup image, as shown below:

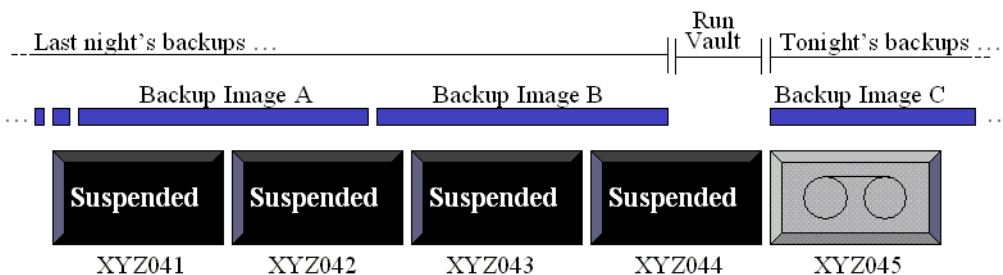
Without Suspend



To avoid vaulting (ejecting) partial images, use one of the three following methods:

- ◆ Stop backup activity long enough to run Vault.
- ◆ Vault backups which are older than one day.
- ◆ Suspend all tapes on which backups have been written within the last day.

Using the Suspend Option



Suspend is a useful option in several situations:

- ◆ When Vault is configured to eject tapes which are created by something other than the Vault process (usually the original NetBackup backup process).
- ◆ When backups are scheduled around-the-clock so that there is no convenient pause in backup activity to allow Vault to operate. For example, if Vault is to eject the tapes created by original backup jobs, then you would configure the Suspend option to suspend all the tapes from today that you want to vault tomorrow. This way, no more backups will be written to those tapes and they will be ready to be vaulted. (They will not contain any partial backups).

Note Vault only suspends the tapes in off-site volume pools specified in the Eject off-site volume pool list.

Vaulting Original Backups in a 24x7 Environment

To vault up-to-the-minute original backups in a 24x7 backup environment, the user may want to vault tapes while backups are currently being written. The user does not want to pause backups long enough to suspend tapes.

You cannot suspend tapes on which backups are currently being written. Therefore, partial images will be vaulted.

Vault may attempt to eject a busy tape. This will cause an error. The rest of the image will be vaulted next time the Profile is run.

Note If Vaulting originals, wait a day. Choose backups that were queued a day or more ago. (This assumes that your backups will be complete by the time the Vault session runs.)

Preparing for Efficient Recovery

Revault Media

If the primary copy of a backup is recalled from a piece of vaulted media, and you do not revault that media, then none of the images on that media -- even primary images -- will be recoverable.



Use Good Naming Conventions for Volume Pools

For volume pools that have tapes that are (or need to be) vaulted, choose a volume pool name that makes it easy for others to recognize it as a Vault volume pool. Good volume pool names would be Vaulted_Payroll and Vaulted_CustomerDB or even 1_month_vault and 7_year_vault. This naming convention will help you (and others) to organize and more easily identify media in the future.

Specify a Unique Volume Pool For Each Logical Vault

Specify a unique volume pool for each logical vault you configure. Duplicated images are assigned to a volume pool. We recommend that images intended for off-site vaulting be assigned to a volume pool specified for this purpose. Once a tape is assigned to a pool it remains there and will be used for rotation within that same vault.

Specifying the Primary Copy and Keeping It On Site

When vaulting a copy of a backup, keep your primary copy in your on-site robotic library.

When vaulting a copy of a backup, you will want to ensure that the copy which remains in your on-site robotic library is the primary copy. This is because NetBackup will always use the primary copy for restores. (If the offsite copy were primary, then when the user tried to restore a file, NetBackup would wait for a mount of the off-site media.)

If you are using Multiple Copies in your backup policy to create the on-site and the off-site copy during the backup process, make sure that the first copy on the Multiple Copies dialog is the on-site copy. During the backup process, the first copy is always the primary copy.

If your Vault profile is using duplication to create the on-site copy, be sure to configure Vault to make the on-site copy primary. (Some sites prefer to vault the original backup because it is theoretically more reliable than a duplicated copy.)

Avoid Resource Contention During Duplication

Note If you vault only original backups, you need not concern yourself with this topic.

Certain best practices help you avoid resource contention. When we talk about resources, we mean any of the following:

- ◆ Time
- ◆ Media

- ◆ Robots and drives
- ◆ Bandwidth

When Two Processes Try to Use the Same Drive

One problem that can occur as a result of naive configuration is deadlock. This can happen, for example, if two processes try to use the same drive at the same time.

To avoid this deadlock situation, follow the advice provided in either Alternative A or Alternative B below. These alternative configurations work well for multi-robot environments. They make good use of available resources and are unlikely to cause resource allocation problems.

Alternative A: Single Dedicated Robot for Vault Processing

In a multi-robot environment, dedicate one robot strictly for vault processing. The media in this robot will contain only the duplicate backup copies that are to be ejected and sent to the off-site vault. This configuration works best in a SAN environment where all media servers have direct access to the vault robot, because then the duplication step will not send data over the network.

If you choose this configuration, do not use *Any_Available* storage unit in your backup policies unless *only* your Vault storage units are set to On-Demand Only. Using *Any_Available* for other storage units could cause non-vault backups to be written to the vault robot. You can achieve the same behavior provided by *Any_Available* storage unit by configuring your backup policy to use a storage unit group that includes all storage units *except* for the vault robot's.

There are two ways of achieving this configuration:

- ◆ Use the multiple copies (inline tape copy) feature in your backup policies. Send the first (the primary) copy to a storage unit group that is not in the vault robot. Send another copy (the second or third or fourth, depending upon how many copies you are making) to the vault robot and to the off-site volume pool. Configure a Vault profile to eject all media in the vault volume pool. This configuration requires that your media server be connected to both robots.
- ◆ Use Vault to do duplication. Backup images will be duplicated from all other robots to the vault robot. If you want to avoid sending duplication data over the network, you must take great care in choosing which media server(s) will do the duplication. See “Avoid Sending Duplicates Over The Network” beginning on page 40.) The following alternatives are available via the **Duplication** tab of the **Change Profile** dialog in the GUI:



- Do not use the **Advanced** view. Do *not* specify a alternate read server. For each backup image, the media server that did the backup will also do the duplication. All media servers will send duplication data (over the network) to the destination storage unit's media server. (Exception: Backups which were originally written by the media server of the vault profile's destination storage unit will not be sent over the network.)
- Do not use the **Advanced** view, but *do* specify the destination storage unit's media server as the alternate read server. If the alternate read server also has access to all of the backup robots, then no data will be sent over the network.
- On the **Choose Backups** tab, in the **Media Servers** list, specify all media servers. On the **Duplication** tab, use the **Advanced** view. Select the **Alternate read server: Read original....**checkbox. Create an entry in the list for each of the media servers specified on the **Choose Backups** tab.

Note To avoid sending duplication data over the network, make sure that each media server listed on the **Choose Backups** tab is represented as a media server on the advanced view of the **Duplication** tab.

For each entry in the list, you can avoid sending duplication data over the network by specifying the destination storage unit's media server as the alternate read server. The alternate read server must also have access to all the robots which hold the media to be duplicated.

Note If the alternate read server does not have access to the robot that is the source for the duplication process, then the images in that robot will not be duplicated.

Ensure that the total number of write drives specified in the **Write Drives** column of the list does not exceed the number of drives in the vault robot.

Advantage

This configuration is most convenient for the operator, who can eject and inject tapes from only one robot, simplifying the tape rotation process.

Disadvantage

In a complex environment, this configuration can be difficult to configure if you need to avoid sending duplication data over the network.

Alternative B: Configuring Each Robot as a Vault Robot

In a multi-robot environment, configure each backup robot to be a vault robot as well. Each robot will duplicate and/or eject only backup images that were originally written to it. There are three alternative ways of doing this:

- ◆ Let your backup jobs create the copies to be vaulted (by using the Multiple Copies feature in your backup policies). Send all copies to one robot, putting the copy to be vaulted into your off-site volume pool. Only backups on media in the off-site volume group (specified on the Eject tab) and which meet the rest of the criteria specified in the vault profile will be ejected.
- ◆ Use Vault to do duplication: On the **Choose Backups** tab of the Change Profile dialog, specify the **Source Volume Group** to be the robot to which this profile belongs. This will limit the profile so that it will duplicate only backup images that have their primary copy on media in this robot. Take care to specify no more than half of the available drives in the robot as read drives, because just as many drives will be required as write drives.

Advantages

This configuration will easily avoid a deadlock situation that can happen when one profile attempts to do duplication of images originally stored in multiple robots. This deadlock situation can otherwise be avoided, but it requires that only half of the available drives are specified as write drives.

It is easy to avoid sending duplication over the network with this configuration. For the Duplication step, it will not be necessary to use the Advanced view. Do specify the alternate read server to be the media server of the specified destination storage unit.

Note The destination storage unit must have at least two drives if that robot will be used for both read and write functions.

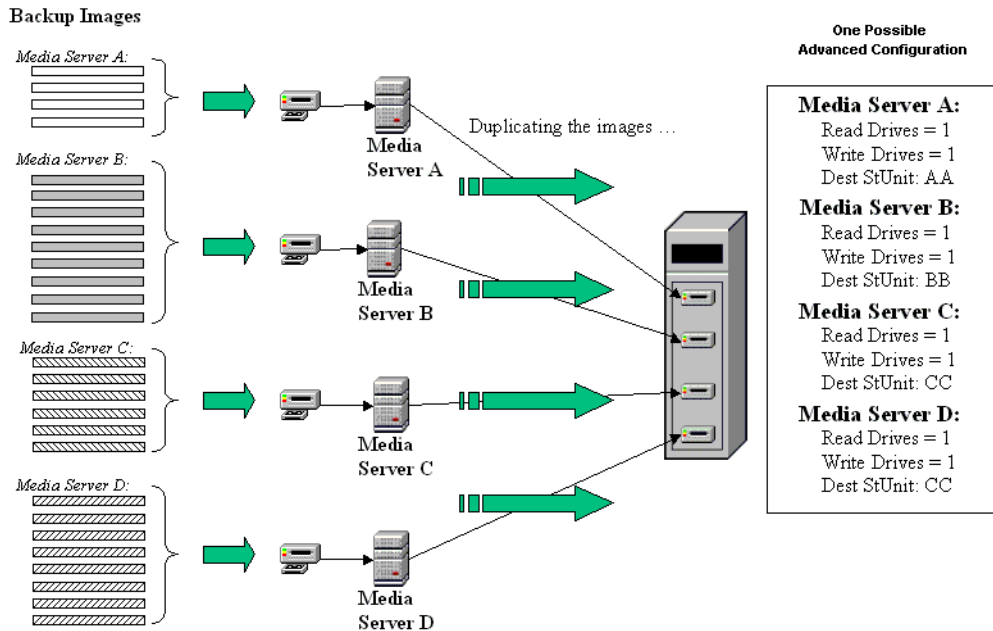
When the Read Drive Is Not in the Vault's Robot

Note If you vault only original backups, you need not concern yourself with this topic.

As the illustration below shows, the read drive need not be in the Vault's robot. For configurations which involve multiple media servers and multiple robots, we recommend that you seek advice from VERITAS Enterprise Consulting Services:



One Profile Using Advanced Duplication



For the above configuration, do not use *Any_Available* storage unit for backups. For the setup shown above, we recommend that you contact VERITAS Enterprise Consulting Services.

Sharing Resources with Backup Jobs

Note If you vault only original backups, you need not concern yourself with this topic.

Obviously, NetBackup and Vault use the same resources, specifically media servers and tape drives. *Any_Available* storage unit would send some original backup images to the vault robot (the robot to the right of the diagram). Subsequently, when Vault tried to duplicate those images, it would need a read drive and a write drive in the vault robot. If not enough drives were available, this could result in a deadlock situation. The easiest way to prevent a potential conflict is to be sure that backup jobs and vault jobs are not scheduled at the same time. Wait until NetBackup has completed its backup jobs before starting a vault job.

This is especially important because Vault 4.5 contains load-balancing logic that requires one vault session to use all configured read and write drives until the last of the tapes is being duplicated. While this is an efficient use of resources, it means that it is more likely that NetBackup and Vault could try to use the same drive if backup jobs and vault jobs are scheduled for the same time period.

We recommend that you review a snapshot of the images you want to duplicate before you run the vault job. This exercise will give you a sense of where the images are located, and what kind of resources will be required. You can create a snapshot using the Preview method discussed in the “Administrative Tasks in Vault” chapter.

Load Balancing

Note If you vault only original backups, you need not concern yourself with this topic.

If it is feasible, we strongly recommend you use the multiple copies (Inline Tape Copy) feature in your backup policies to create both the on-site copy and the copy that will be sent to the vault, rather than using Vault duplication to create either of these copies. Avoiding the Duplication step in a Vault profile avoids all resource contention issues and can significantly simplify the vaulting process.

If Your Vault Vendor Does Not Pick Up Media Every Day

Note If you vault only original backups, you need not concern yourself with this topic.

You can use Vault to spread the duplication workload evenly throughout the week rather than doing all of the duplication on one day and leave the duplicated media in the robot until it is due to be collected by the vault vendor. For instance, suppose the vault vendor is due to pick up the media every Friday. Then:

- ◆ Configure one Vault profile to do only the Duplication step.
 - Configure a Vault policy to schedule this profile to run every day of the week.
 - Configure a `vlt_end_notify.robot.vault.profile` script to run after every execution of this particular vault profile. For more information on notification scripts, please see “Using Notify Scripts.”
 - Program this script to invoke a normal, on-site Catalog Backup (using `bpbackupdb`).
- ◆ Configure a second Vault profile to do the Vault Catalog Backup step and the Eject step.



- Configure this profile to use the same criteria to select the backup images as configured in the previous Vault profile. (In the GUI, this criteria is specified on the **Choose Backups** tab of the **Change Profile** dialog)
- Configure a Vault policy to schedule this profile to run before the vault vendor arrives on Friday.
- Make sure that you put the correct off-site volume pools in the Eject step. Tapes in these pools are the only ones that will be ejected.

Specifying Different Volume Pools for Source and Destination

Note If you vault only original backups, you need not concern yourself with this topic.

You should never configure a profile for duplication such that the source tape and the destination tape are in the same volume pool. This will result in deadlock when NetBackup chooses the same tape as the source and the destination of the duplication operation. (This is a NetBackup limitation.)

Avoid Sending Duplicates Over The Network

Sending duplicates over the network is not a problem if there is sufficient bandwidth, but even a fiber optic SAN can only handle a couple of duplication jobs at a time.

Here are some strategies to avoid sending data over the network.

Use ITC (Multiple Copies) To Minimize Data Transfer

One way to avoid sending data over the network with your Vault job is to use Inline Tape Copy (Multiple Copies) during nightly backups. This avoids the need for your Vault session to do duplication. In this scenario, Vault need only eject the backup tapes. Vault takes no significant resource time, except for the Catalog Backup step. (Catalog Backup is necessary to capture the changed volume database information for each vaulted tape.)

Suppose you want the on-site copy of your backups to go to one robot, and the off-site copy to go to another robot. Inline Tape Copy requires that all destination storage units be on the same media server. Therefore, your media server will need a storage unit on both robots (one storage unit for your on-site copy and one for the off-site copy.)

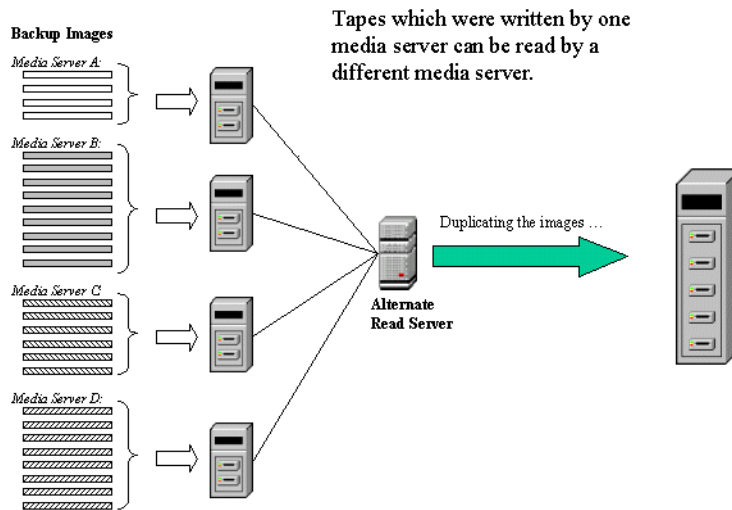
Use Alternate Read Server

Note If you vault only original backups, you need not concern yourself with this topic.

If your vault profile does duplication, then you can avoid sending data over the network by specifying an alternate read server. An alternate read server is a server used to read a backup image originally written by a different media server. It must have access to the robots holding the media with the original backups. To avoid the network, the alternate read server must also be the destination storage unit's media server.

Note If the destination storage unit is not connected to the alternate read server, you will send data over the network.

For example, in the diagram below, non-disk images written by media servers A, B, C, and D will be read by the alternate read server.



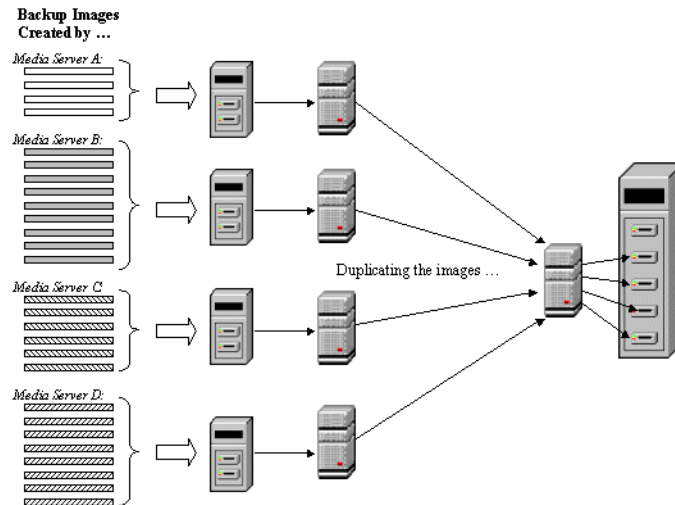
If the destination storage unit also belongs to the alternate read server, then no data will go over the network.

Use Advanced Duplication Configuration

Note If you vault only original backups, you need not concern yourself with this topic.



If each media server has access to at least one unique drive in the destination robot, you can use advanced duplication configuration to process each media server independently and concurrently. (Note: all media from a single profile are ejected from the same robot.)



With the same connectivity as in the above diagram, in theory you could do the same thing by configuring a separate profile for each media server, rather than using advanced duplication configuration. However, multiple profiles within a single vault must run serially; so this may not allow you sufficient bandwidth.

Taking Care When Specifying All Media Servers

Note If you vault only original backups, you need not concern yourself with this topic.

Be sure not to specify **All media servers** on the Choose Backups tab of a profile if you are planning to use the Advanced Configuration option on the Duplication tab.

If you list more media servers on the Choose Backups tab than on the Duplication tab, Vault assigns the images written by media servers not listed in the Advanced view to the first media server that finishes its duplication job. If the first available media server is across the network, a lot of data would be sent over the network.

Another possible, though less problematic, consequence is that backup images from the media servers not configured for duplication may be duplicated by a different media server each time the profile is run.

Increase Duplication Throughput

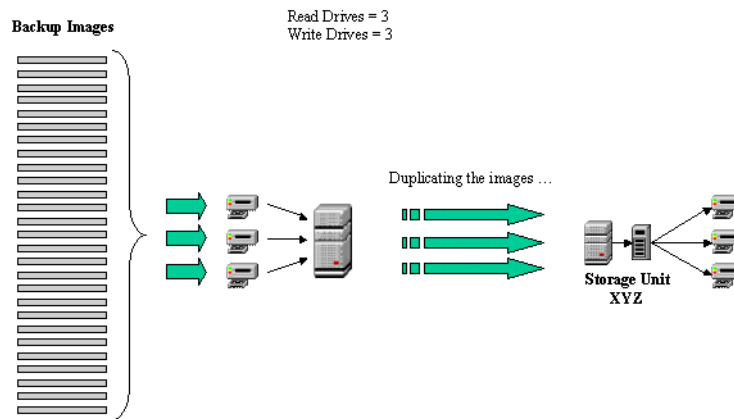
Note If you vault only original backups, you need not concern yourself with this topic.

Adding drives will enable Vault to run multiple `bpduplicate` sessions concurrently. For each Write Drive, a separate `bpduplicate` job will be forked.

Configuring for Multiple-Drives: Basics

Note If you vault only original backups, you need not concern yourself with this topic.

In the scenario below, there are three read drives and three write drives. Three `bpduplicate` processes will run concurrently.

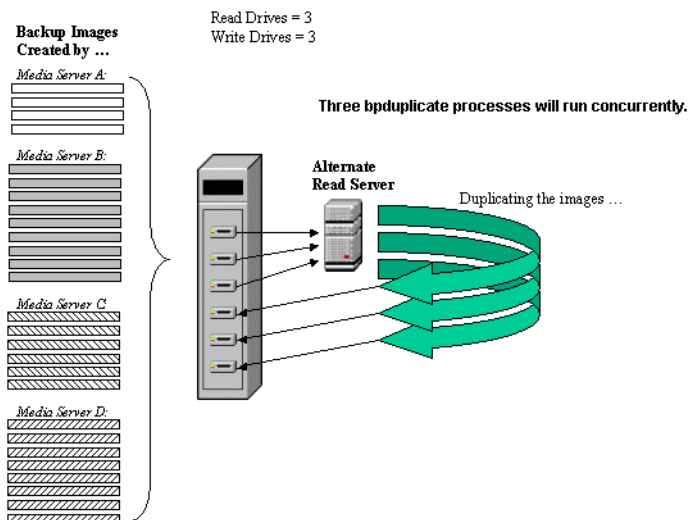


Multiple-Drive Scenario: Does Not Send Data Over Network

Note If you vault only original backups, you need not concern yourself with this topic.

The alternative scenario shown below uses three read drives and three write drives, like the one above, but it avoids sending data over the network.





In this arrangement, three bpduplicate processes are running concurrently. The non-disk images written by media servers A, B, C, and D will be read by media server A. If the destination storage unit also resides on the alternate read server (media server A), then no data needs to be sent over the network.

Be sure you have a thorough understanding of your NetBackup configuration and of your company's vaulting procedures before you begin to configure Vault.

This chapter discusses the following:

- ◆ Methods of Configuration
- ◆ Configuring Robots for Vault
- ◆ Creating a Vault
- ◆ Creating a Profile
- ◆ Configuring a Profile

Be sure you have a thorough understanding of your NetBackup configuration and of your company's vaulting procedures before you begin to configure Vault.

Methods of Configuration

Vault configuration can be done using any of three interfaces:

- ◆ Java Graphical User Interface (GUI)
- ◆ NT GUI
- ◆ `vltadm` -- a Menu User Interface (MUI) analogous to `bpadm`.

When to Use `vltadm`

Although this chapter presents the configuration of Vault from a GUI perspective, the same configurations can be accomplished using `vltadm`.

- ◆ You must use `vltadm` if:
 - You are running on a UNIX box that does not support NetBackup Vault's Java GUI, and there is no Windows NetBackup Console from which to remotely manage the server. For example, this situation would exist for SGI and Sequent platforms.



- You find it necessary to manage NetBackup by dialing up from elsewhere without a GUI. For example, if the Administrator's beeper goes off in the middle of the night, it is necessary to be able to resolve problems using MUIs.
- ◆ You may want to use `vltaadm` if:
 - You have servers in the DataCenter with a terminal only and want to use that machine, rather than going elsewhere to connect remotely from a desktop.
 - You hate GUIs.

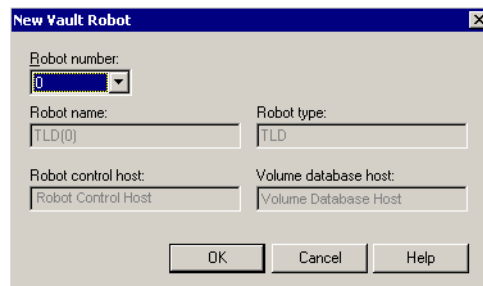
Configuring Robots for Vault

Your first step in using Vault is to choose Vault robots. These are the robots that will hold the media with original or duplicated images, and from which media to be vaulted will be ejected. These robots must already be configured through NetBackup.

▼ To tell Vault what robots to use:

1. Highlight **Vault Management**.
2. Open the **Actions** menu and select **New** and then **New Vault Robot**.

The New Vault Robot dialog displays.

The image shows a Windows-style dialog box titled "New Vault Robot". It contains several input fields: "Robot number:" with a dropdown menu showing "0"; "Robot name:" with a text field containing "TLD(0)"; "Robot type:" with a text field containing "TLD"; "Robot control host:" with a text field containing "Robot Control Host"; and "Volume database host:" with a text field containing "Volume Database Host". At the bottom, there are three buttons: "OK", "Cancel", and "Help".

Based on the robot number that you select, the other fields will be filled in automatically. Eligible robots are recognized by Vault.

3. Select a robot number. NetBackup assigns a number to each robot that it recognizes. The rest of the dialog will be completed for you.

Creating a Vault

You must define at least one logical vault for each robot. A robot may contain multiple vaults, but a vault cannot span robots. Each vault is associated with one volume group.

Never use the same Off-site Volume Group for more than one logical vault. Otherwise individual reports will pick up tapes from multiple logical vaults.

You must specify a name and a physical location (on a storage device) for the vault. For example, if your NetBackup Vault configuration contains three distinct TLD robots (not connected with pass-through devices), you would define at least three logical vaults, one for each TLD robot.

Vault automatically creates a directory for each vault in the following directories:

- ◆ UNIX: `/usr/opensv/netbackup/vault/sessions/vault_name`
- ◆ Windows: `install_path\netbackup\vault\sessions\vault_name`

These directories are created the first time a session is run for this vault. They will hold information about each vault session.

▼ To create a vault:

1. Expand **Vault Management** on the console tree.
2. Highlight a robot.
3. From the **Actions** menu, choose **New** and then **New Vault**.

The **New Vault** dialog displays.

4. Enter a vault name.

The name of the vault should reflect its purpose. For example, if you are creating this vault primarily to duplicate records from the finance department, you might call the vault `Finance_Dups`. The vault name may contain up to 25 characters.



5. Select the robotic volume group for the vault.

The robotic volume group is the volume group from which media to be vaulted will be ejected.

6. Enter the first off-site slot ID.

Off-site slot IDs are usually used by the vault vendor to track media. If your vendor does not use these identifiers, you can use the default first off-site slot ID of 1. Off-site slot IDs are unique only within a given vault. The first off-site slot ID is sometimes referred to as the off-site slot seed.

Slot IDs are assigned contiguously from the starting slot number. Ensure that the number of media in the vault does not exceed the range of slot IDs assigned by the vault vendor. With every session, Vault starts with the off-site slot ID and counts upwards, looking for slots that are no longer in use. Vault always fills in the gaps with newly vaulted media.

In case multiple vaults are defined for the same vault vendor, you must divide the range of assigned slots between the various vaults. For example, if the vault vendor has assigned the range 1-2000 and you have defined 3 vaults for this vault vendor, then you can assign range 1-499 to vault 1, 500-999 to vault 2, and 1000-2000 to vault 3, assuming vault 3 has the maximum number of tapes to vault.

7. Identify the vault vendor.

8. Enter the off-site volume group for which you are creating the vault.

The off-site volume group should have a unique identifier which might be the same as the vendor name or location. Vault moves each piece of media to be ejected from the robot into a non-robotically-controlled volume group. This step ensures that Media Manager does not try to add more images onto the same tape and tells NetBackup that the volume is not available for use. The off-site volume group name may contain up to 25 characters.

Creating a Vault Policy

Setting up a Vault policy differs from setting up a regular policy in NetBackup. Be sure to specify Vault as the policy type, and rather than entering a directive on the File tab, indicate one of two Vault commands. There are no clients specified in Vault policies.

▼ **To create a Vault policy:**

1. In the NetBackup Administration window, expand **Master Server > NetBackup Management > Policies**.
2. Click the **New Policy** button.
3. Type a unique name for the new policy in the **Add a New Policy** dialog.
4. Click **OK**.
5. On the Attributes tab, select **Vault** as the policy type.
6. On the Schedules tab, click **New** to create a new schedule.
The type of backup defaults to **Automatic**.
7. Complete the schedule.
8. Bypass the Client tab, since Clients are not specified for Vault jobs.
9. On the Files tab, enter one of two Vault commands. Both commands are found in:
UNIX: `/usr/opensv/netbackup/bin`
Windows: `install_path\NetBackup\bin`
 - `vltrun`: Use to specify the robot, vault name, and profile for the job. Enter:
`vltrun profile_name`
If the profile name is not unique, enter:
`vltrun robot_number/vault_name/profile_name`
 - `vlteject`: Use to eject media and/or generate reports for previously run sessions. For example:
`vlteject -eject -report [-vault vault_name [-sessionid id]] [-auto y|n] [-eject_delay seconds]`
10. Click **OK**.



Creating a Profile

You must have at least one profile for each vault.

Profiles are templates for vault jobs. They provide an organized way for you to specify which steps you want a vault session to execute.

Specifically, you can configure the following steps:

- ◆ **Configure image selection criteria**

Define the criteria to use for selecting images to be vaulted.

- ◆ **Configure duplication (optional)**

Configure the resources for duplicating the selected images.

- ◆ **Configure catalog backup (optional)**

Define the directory paths for the catalogs to be backed up. You also need to define the volume pool that contains the write media.

- ◆ **Configure eject (optional)**

Identify the pools from which the media is to be ejected. The eject mode (manual vs. automatic) is also configured.

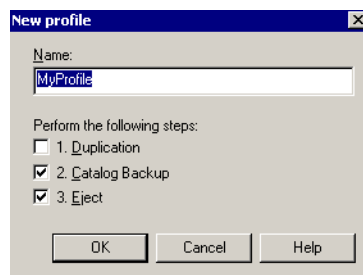
- ◆ **Configure reports (optional)**

Select the various report types and the manner in which the reports are to be dispersed, i.e., via printer, email, or a directory. The report generation mode (manual vs. automatic) is also configured.

▼ To create a profile

1. Highlight a vault on the NetBackup console tree. From the **Actions** menu, choose **New** and then **New Profile**.

The New Profile dialog displays.



2. In the **Name** field, type a name for the profile. We suggest you make it as descriptive as possible.
3. Select the steps you want this profile to perform.

You must select at least one step, although you might choose to skip one or two of these steps, depending on the primary purpose of the profile. As you configure the profile, you have the opportunity to change your mind about the selections you made in this dialog.
4. Click **OK**.

The New Profile: *name_of_profile* dialog displays. See “Configuring a Profile” below for instructions on configuring the tabs of this dialog.

Configuring a Profile

The New Profile... dialog includes the following five tabs:

- ◆ The **Choose Backups** tab is where you specify the criteria for selecting backup images. You can choose from the list of available criteria. The criteria are based on:
 - Time window
 - Types of backup
 - Source volume group
 - Clients
 - Media Servers
 - Policies
 - Schedules
- ◆ The **Duplication** tab is where you configure duplication of the selected backup images. For instance, the number of copies to make and the resources to use (which media servers, storage units, and how many drives) are specified here.
- ◆ The **Catalog Backup** tab is where you choose how to back up the NetBackup and Media Manager catalogs. You must vault a new catalog backup for efficient disaster recovery.
- ◆ The **Eject** tab is where you choose in which off-site volume pools Vault should look for the media you want to eject.
- ◆ The **Reports** tab is where you choose what reports to generate.



A profile can include any or all of these options. The following pages will outline the procedures for image duplication, catalog backup, and media eject. For information about configuring the reports tab, and for consolidating eject and report actions, please see “Reporting” beginning on page 109.

Profiles for Both Originals and Duplicates

You can create a profile that vaults both originals and duplicates. For example:

- ◆ NetBackup policy A creates multiple copies and assigns one of these copies to the off-site volume group.
- ◆ NetBackup policy B creates one copy and assigns it to an on-site volume group.
- ◆ Your Vault profile is configured to duplicate backups to the off-site volume group.

When you run a Vault session, Vault will only duplicate backups from NetBackup policy B because it will know that there is already an original backup from policy A in the off-site volume group. If you have configured the profile for eject, it will eject both duplicates (from policy B) and originals (from policy A).

Configuring the Choose Backups Tab

The first step in creating a profile is to identify the criteria by which Vault selects backups. For doing so, you will use the Choose Backups tab, shown below:

The screenshot shows the 'New Profile...: dec4' dialog box with the 'Choose Backups' tab selected. The dialog has a title bar with a close button. Below the title bar is a tabbed interface with four tabs: '1: Duplication', '2: Catalog Backup', '3: Eject', and 'Reports'. The '1: Duplication' tab is active. Inside this tab, there is a text box with the instruction 'Please choose the criteria for selecting the backups which you would like to vault.' Below this, there are two rows of time selection controls. The first row is labeled 'Backups started:' and has 'between' followed by a spinner set to '8', 'day(s)', and a spinner set to '0', 'hour(s) ago'. The second row is labeled 'and' followed by a spinner set to '1', 'day(s)', and a spinner set to '0', 'hour(s) ago'. Below these is the text 'relative to start time of the session.' To the right of these controls is a checkbox labeled 'Enable advanced criteria' which is currently unchecked. Below the checkbox are two dropdown menus: 'Type of backup:' with '<All backups>' selected, and 'Source volume group:' with '<All volume groups>' selected. At the bottom of the dialog, there are four columns of selection controls: 'Clients:' with '<All clients>', 'Media servers:' with '<All media servers>', 'Backup policies:' with '<All backup policies>', and 'Schedules:' with '<All schedules>'. Each column has a 'Change...' button below it. At the very bottom of the dialog are 'OK', 'Cancel', and 'Help' buttons.

The most basic criterion you can set is the time frame. Select the **Enable Advanced Criteria** box to make more detailed selections.

▼ **To specify criteria for selecting backup images**

1. Indicate a time range (in terms of days and hours) for the backup images you want to vault.

Time is expressed in terms of days and hours, relative to the time of the session. For example, assume the setting is as below:

between 8 day(s) 0 hour(s) ago

and 1 day(s) 0 hour(s) ago

If the session is started on October 12 at 1:00 pm, count backwards from October 12. The vaulted backups will be those started between October 4 at 1:00 pm (8 days before) and October 11 at 1:00 pm (1 day before).

2. If you want to specify any other details on this screen, check **Enable advanced criteria**. By default, **All** is selected for the categories on this screen.

- a. Specify the types of backup to be vaulted.

Depending on the different types of backups you have configured in Policy Management, you can choose the backup type as a criterion. Only those types for which you have configured policies will be available for selection. This field is optional. If you want to vault all types of backups, accept the default.

- b. Specify the source volume group.

From the dropdown list, select the source volume group, which is the robotic volume group from which backups can be selected. This is mainly useful when duplicating backups in a multiple-robot configuration if you wish to only select backups based on the physical location of the media.

Note If this profile skips the Duplication step, but executes the Eject step, only backups that already belong to the robotic volume group can be ejected. If you specify a source volume group that is different from the robotic volume group in this situation, no tapes will be ejected because none of the backups in the source volume group are also in the robotic volume group unless they have been duplicated.

- c. In the Clients list box, to change the default, **All**, click the **Change** button. The Clients dialog box displays. Choose the clients you want to include in this profile. To return to the default, select the **Include all...** checkbox in the upper left of the dialog box.



- d. In the Media Servers list box, to change the default, **All**, click the **Change** button. The Media Servers dialog box displays. Choose the media servers you want to include in this profile. To return to the default, select the **Include all...** checkbox in the upper left of the dialog box.

For every media server you select, make sure there is an entry on the Alternate Media Server Names tab of the Properties dialog. At a minimum, there should be, for each media server, an entry which contains both the abbreviated name and the fully qualified name. Also add any other names by which a media server has ever been known. Taking this action will avoid a number of problems. For example, if you do not list alternate names for media servers, some images may not be recognized to match the Choose Backups criteria and may therefore not get vaulted.

- e. In the Backup Policies list box, to change the default, **All**, click the **Change** button. The Backup Policies dialog box displays. Choose the Backup Policies you want to include in this profile. To return to the default, select the **Include all...** checkbox in the upper left of the dialog box.

If you select the **Exclude** checkbox in the upper right of the dialog box, the policies you have selected will be excluded from the profile, and all others will be included. Use this option if there are only a few policies you do not want to include in the profile.

- f. In the Schedules list box, to change the default, **All**, click the **Change** button. The Schedules dialog box displays. Choose the schedules you want to include in this profile. To return to the default, select the **Include all...** checkbox in the upper left of the dialog box.

- 3. When you have chosen all your settings, move to the next tab.

Vault compares images in the NetBackup database with these criteria and generates a list of those that qualify. This process, which is called the Image Selection process, creates two output files called `image.list` and `preview.list`, which are placed in the following directory:

UNIX: `/usr/openv/netbackup/vault/sessions/vault_name/sidxxx`

Windows: `install_path\NetBackup\Vault\sessions\vault_name\sidxxx`

where `vault_name` is the name of the vault and `xxx` is the session identifier.

Configuring the Duplication Tab

Note If you are only vaulting originals, you need not concern yourself with the complexities of duplication rules.

In Vault, configuring duplication is predicated upon the configuration of duplication rules. A duplication rule describes how an image will be duplicated. It includes the following parameters:

- ◆ Source Media Type (Disk and/or Removable?)
- ◆ Alternate Read Server (Optional)
- ◆ Number of Read Drives
- ◆ Destination Storage Unit (may be up to 4 of these)
- ◆ Number of Write Drives
- ◆ Destination Volume Pool (may be up to 4 of these)

If only one duplication rule is specified, then all backups images are copied according to the same rule. This is what we see when we look at the standard (not the Advanced) view of the Duplication tab. In this case, the source media server is not specified, and therefore all selected backups that include the specified media servers are duplicated.

Creating Duplication Rules of Different Complexity

You can use the basic Duplication tab to do the following tasks, listed below in order of their level of complexity:

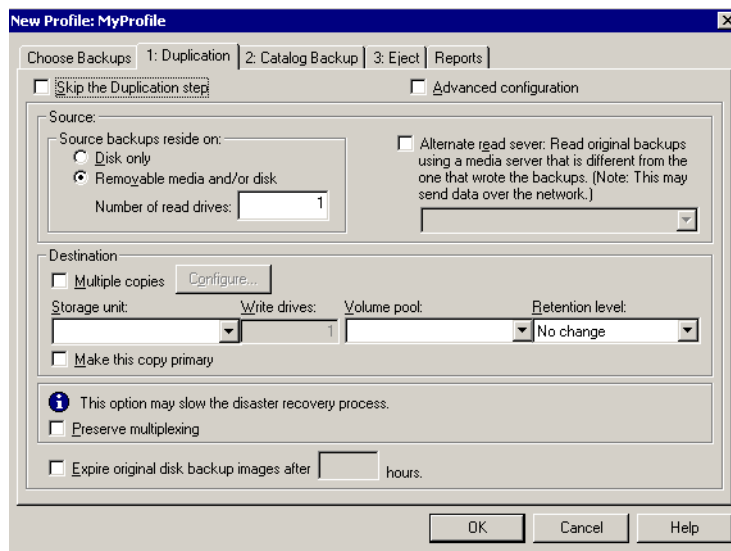
- ◆ Enter basic information instructions for making a single duplicate of an image
- ◆ Access the Multiple Copies dialog
- ◆ Access the Advanced Configuration options
- ◆ Access the Duplication Rule dialog

For more information on duplicating backup images, see the *NetBackup System Administrator's Guide, Duplicating Backup Images*.



Basic Duplication

The screen shot below shows the basic Duplication tab:



Note For every media server, add an entry on the Alternate Media Server Names tab of the Properties dialog. At a minimum, there should be, for each media server, an entry which contains both the abbreviated name and the fully qualified name. Also add any other names by which a media server has ever been known. Taking this action will avoid a number of problems. For example, if you do not list alternate names for media servers, some images may not be recognized to match the Choose Backups criteria and may therefore not get vaulted. If you have multiple NIC cards in your server, make sure that the server name or IP address associated with each NIC card is listed in the Alternate Media Server Names dialog. See “Adding Alternate Media Server Names” on page 95 for more information.

▼ **To complete the basic Duplication tab:**

1. Familiarize yourself with the different subscreens you can access from the Duplication tab:
 - a. If you select the **Advanced configuration** checkbox, the screen will change to allow you to specify different options. If you take this step, see “Duplication: Completing the Advanced Configuration Dialog.” Otherwise, continue with step 2 below.
 - b. If you select the **Multiple copies** checkbox, you must click **Configure** to bring up the Multiple Copies dialog. See “Duplication: Completing the Multiple Copies Dialog” for information on completing that dialog.
2. Indicate whether the images you want to duplicate reside on disk storage units only or on disk and/or media storage units.
3. Enter the number of drives to be used for reading backup images for duplication.

When you enter a number of read drives, the same number will be entered into the destination Write Drives field. You must have an equivalent number of read and write drives available.
4. If robots (or drives if you are using SSO) are shared by more than one media server, you can designate a different media server to read the original backups than the media server that wrote the backups. By default this option is turned off. If you want to turn on this option, check the **Read original backups...** box and enter the name of the media server you want to use for reading.
5. Select the **Multiple copies** checkbox if you want to make more than one copy.

Note If you do select **Multiple Copies**, you must click the **Configure** button to continue. See “Duplication: Completing the Multiple Copies Dialog” for information on completing that dialog. Otherwise, continue with step 6 below.

6. From the dropdown list, select a storage unit for the duplicate.
7. Select the off-site volume pool where Vault will place the duplicate image. *Do not* use the volume pool which was used for the original backup, because this may result in deadlock when two processes try to use the same volume pool at the same time.
8. From the drop-down list, identify the retention level for this copy.



Each image copy has a separate expiration date. When a duplicate is created, its expiration date will be the same as the original, if a retention level is not specified in the copy's configuration.

9. Select the **Make this copy primary** checkbox if you want this to be the primary copy.

If you do not select the copy as primary, then the original will be the primary copy. The primary copy is the one from which restores are made, and it should be kept on site.

10. Indicate whether you want to preserve multiplexing.

Multiplexing is the process of sending concurrent-multiple backup images from one or more clients to the same piece of media. This process speeds up duplication, but slows down restores and disaster recovery processes. If the option to preserve multiplexing is selected, the multiplexed duplication process will occur for all multiplexed images that are selected for duplication during a given Vault session. In bpvault 3.4, the default was to demultiplex all images.

Note If multiplexing is enabled in Vault, make sure that all destination storage units have multiplexing enabled. (Set through Media Manager -- **Global Configuration of Storage Unit.**)

11. Check **Expire original disk backup images after *nnnn* hours** and specify the number of hours from the time the Vault session runs to expire the original backup.

You can use this option to force an earlier expiration time for the images so the disk space is freed up for use by subsequent backups. If the duplication of a disk image is not successful, the disk image will not be expired.

Note Consider expired backups inaccessible. Only configure Vault to expire a disk backup if you are sure you will not need it.

12. When you have filled in the appropriate information, click **OK**.

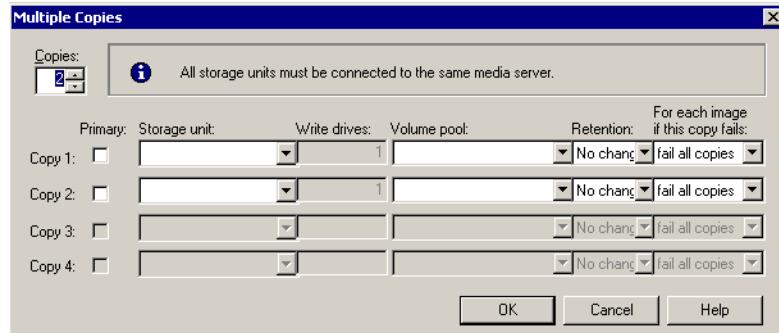
Duplication: Completing the Multiple Copies Dialog

The Multiple Copies screen becomes available only if you have selected the **Multiple Copies** checkbox on the Duplication tab screen and have clicked **Configure**. You can return to basic duplication by clicking **Cancel** and deselecting **Multiple Copies**.

▼ **To complete the Multiple Copies option:**

1. In the **Copies** spinbox, indicate how many copies to make.
By default, you will create a single duplicate, but you can make up to 4.

Note All storage units must be connected to the same media server.



2. For each copy, specify the storage unit, volume pool and retention level, and indicate what action is to be taken if the copy fails. We suggest that you assign only one copy of the backup to the off-site volume pool before you duplicate the backup.
3. Select the checkbox next to the copy you want to designate as the primary copy, if you want to change the primary copy.

The primary copy is the copy from which restores are generated. It should be kept on site.

4. Click **OK**.

For more information about multiple copies, see “Inline Tape Copy.”

Duplication: Completing the Advanced Configuration Dialog

The Advanced Configuration option is available only if you selected the **Advanced configuration** checkbox on the Duplication tab screen.

Note You never need to use the Advanced view of the Duplication tab if your profile is duplicating images backed up by a single media server. To return to the basic duplication view, deselect the **Advanced Configuration** checkbox.

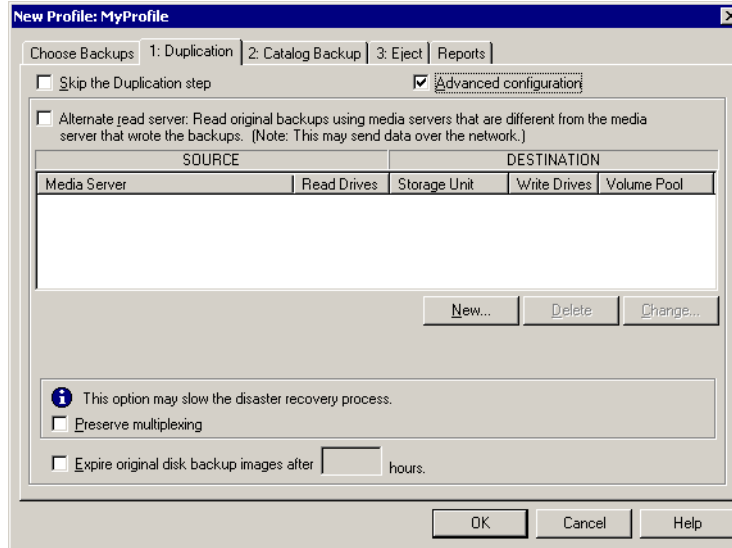


Use the Advanced view of the Duplication tab only if you need to control exactly how to allocate the backup images to be duplicated. The following scenarios may make this concept clearer:

- ◆ Your robot might have different types of media/drives so that you have different storage units to use as destinations of the duplication process. In this case, you may want to spread the load between the two storage units. For instance you may want to send the duplicate copies of all backup images written by one media server to a storage unit of one density and of all backup images written by another media server to a storage unit of another density.
- ◆ Your profile may be duplicating backup images where each media server writing backups was backing up a different type of data, requiring a different retention period. For instance, if media server A backs up your customer database, and media server B backs up warehouse inventory data, then you might want to keep your customer database in your vault for a longer period of time (a different retention) than your warehouse inventory data.
- ◆ You may have one media server that you need to offload from the responsibility of duplicating backup images. For instance, suppose you want to spread the load of duplication across multiple media servers, but there is one media server that does backups but that you do not want to use for doing duplication. Then for that one media server you would specify an alternate read server, and you would let the rest of the media servers handle their own duplication.

Note For every media server, add an entry on the Alternate Media Server Names tab of the Properties dialog. At a minimum, there should be, for each media server, an entry which contains both the abbreviated name and the fully qualified name. Also add any other names by which a media server has ever been known. Taking this action will avoid a number of problems. For example, if you do not list alternate names for media servers, some images may not be recognized to match the Choose Backups criteria and may therefore not get vaulted. If you have multiple NIC cards in your server, make sure that the server name or IP address associated with each NIC card is listed in the Alternate Media Server Names dialog.

The screen shot below shows the Duplication tab when **Advanced configuration** has been selected:



1. If robots (or drives if you are using SSO) are shared by more than one media server, you can ask to read the original backups using a different media server from the one that wrote the backups. By default this option is turned off, because it can send data over the network and can slow performance. If you want to turn on this option, check the **Read original backups...** box.

2. Click **New** and complete the Duplication Rules dialog.

If you selected **Read original backups...**, the Duplication Rules dialog will have fields for both media server and alternate read server. If you did not select the option, the dialog will display a field only for the media server. See “Duplication: Multiple Duplication Rules Option” for instructions on completing the Duplication Rules dialog.

3. Indicate whether you want to preserve multiplexing.

Multiplexing is the process of sending concurrent-multiple backup images from one or more clients to the same piece of media. This process will speed up duplication, but may slow down disaster recovery. The multiplexed duplication process will occur for all multiplexed images that are selected for duplication during a given Vault session.

Note If the source image is multiplexed and the **Preserve Multiplexing** option is selected on the Duplication tab, ensure that the destination storage unit configured for each copy has multiplexing enabled.



4. Check **Expire original...** to select the number of hours after this Vault session completes that disk images will expire. You may use this option to expire disk images to create space for subsequent backup images. If the duplication of a disk image fails, the disk image will not be expired.
5. When you have filled in the appropriate information, click **OK**.

Duplication: Multiple Duplication Rules Option

If you use the Advanced view of the Duplication tab, then you can specify more than one duplication rule. If you do specify more than one duplication rule, Vault needs to know which rule to use for each backup image to be copied.

To identify which rule to use for a backup image, you will specify a unique Media Server for each rule. Then when Vault runs, for each backup image it looks at the Media Server that wrote the backup image, and applies the duplication rule corresponding to that Media Server. In this context, the Media Server does not have any effect other than to identify which rule to apply for a given image.

For example, if an image was written by media server A, then Vault will duplicate the image based on the duplication rule specified for media server A.

To avoid sending data over the network, do the following:

- ◆ For each duplication rule that does not specify an alternate read server, ensure that the Media Server is the same server as the Destination Storage Unit's media server.
- ◆ For each Duplication Rule that does specify an alternate read server, ensure that:
 - The alternate read server is connected to all robots having backup images written by the Media Server corresponding to this rule, and
 - The alternate read server is the same server as the Destination Storage Unit's media server.

If a Duplication Rule does not specify an alternate read server, then the media server that originally wrote the backup image will be used to read the original backup image during the duplication process. Under certain circumstances (see below) this will not be the Media Server associated with the Duplication Rule.

▼ To complete the Duplication Rules dialog:

1. Select the source media server. If the **Read original backups...** option was selected, you must also specify the alternate read server. The alternate read server and source media server may be the same.

Note Please refer to “Best Practices” for information on configuring the alternate read server.

2. Select the number of copies to make.

Note The number of copies you choose on this screen cannot exceed the number of copies specified in the **Maximum Copies** field for the NetBackup master server. (Check **Host Properties/Global NetBackup Attributes**.) By default the value is two, which means one original backup and one copy.

3. For each copy, specify the storage unit, volume pool and retention level, and indicate what action is to be taken if the copy fails. We recommend that you assign only one copy of the backup to the off-site volume pool before you duplicate the backup. See “Using Inline Tape Copy (Multiple Copies) in Vault” for more information about completing this dialog.

4. Select the checkbox next to the copy you want to designate as the primary copy.

If you do not select one of the copies as primary, then the original will be the primary copy. The primary copy is the one from which restores are made, and it should be kept on site.

5. Click **OK** to return to the advanced configuration area of the Duplication tab screen.

6. Indicate whether you want to preserve multiplexing.

Multiplexing is the process of sending concurrent-multiple backup images from one or more clients to the same piece of media. This process will speed up duplication, but may slow down disaster recovery. The multiplexed duplication process will occur for all multiplexed images that are selected for duplication during a given Vault session.

Note If the source image is multiplexed and the **Preserve Multiplexing** option is selected on the Duplication tab, ensure that the destination storage unit configured for each copy has multiplexing enabled.

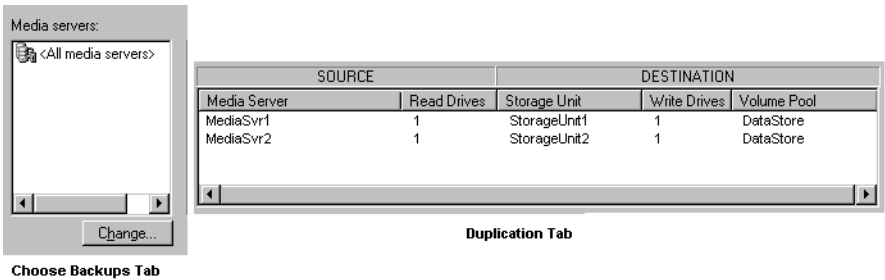
7. Check **Expire original...** to select the number of hours after this Vault session completes that disk images will expire. You may use this option to expire disk images to create space for subsequent backup images. If the duplication of a disk image fails, the disk image will not be expired.

8. When you have filled in the appropriate information, click **OK**.

Treatment of Images Without Corresponding Duplication Rule

In some cases, the profile may list more media servers in the Media Servers list on the **Choose backups** tab (left) than in list on the **Duplication** tab (right):





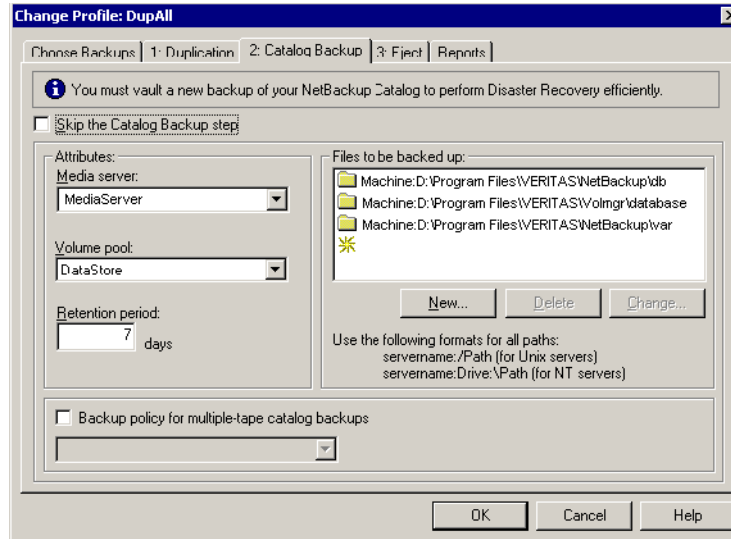
If this happens, then images written by media servers having no corresponding duplication rule must also be duplicated. The intent is to minimize total duplication time by keeping as many drives as possible busily writing data until all images are duplicated. This is handled as follows:

- ◆ All images written by media servers that have a duplication rule, are assigned to the appropriate duplication rule.
- ◆ As soon as one duplication rule has finished processing the images assigned to it, Vault will begin to assign images written by other media servers (media servers that have no rule of their own) to the duplication rule that had finished processing.
- ◆ As other rules complete the duplication of their assigned images, they too will be assigned images written by other media servers that have no rule of their own.
- ◆ Eventually all images written by all media servers listed on the **Choose Backups** tab will be duplicated and the duplication step will be complete. If you have more media servers listed on the **Choose Backups** tab than on the **Duplication** tab, there is only one way to ensure that large amounts of duplication data do not get sent over the network:
 - Every duplication rule must specify an alternate read server. For each duplication rule, the alternate read server must be the same as the media server of the destination storage unit(s).
 - All alternate read servers must be connected to all robots having images written by any media server listed on the Choose Backups tab but not on the Duplication tab.

The above configurations are best suited for a SAN environment where all media servers have visibility to all robots.

Configuring the Catalog Backup Tab

You must have an up-to-date catalog to perform effective disaster recovery. Therefore, you must vault an up-to-date catalog backup every time you vault any backup data. The screen shot below shows the Catalog Backup tab:



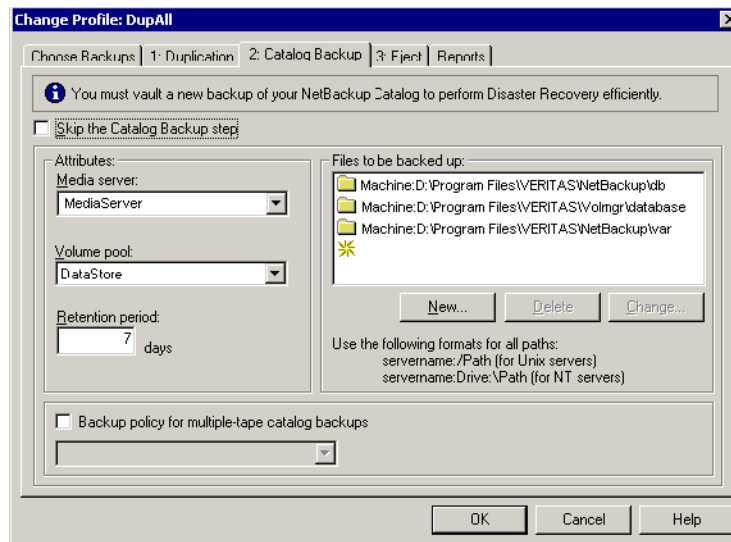
NetBackup requires catalog information to recover backup images. It is extremely important to include catalog backups with your off-site vault rotation. Otherwise it will be much more difficult to recover your business data as you would have to manually import all of your backup tapes to rebuild the catalog. This process takes time that you may not want to spend when you need to resume your business activities.

Vault is capable of backing up the NetBackup catalog to pre-assigned media, ejecting that media and including it in the reports. The catalog media will expire and will be retrieved from the vault after a specified period of time, like other vaulted media.

Note There must be unassigned media in the catalog volume pool before the Vault session runs. If there is no available unassigned media in the pool, the catalog backup will fail.

The NetBackup and Media Manager catalog consists of internal databases that contain information about the NetBackup configuration, any backups that have been performed, the information about backups includes records of the files backed up, and the media on which the files were stored. The catalogs also have information about the media and storage devices that are under the control of Media Manager.





Backing up the Catalog Using Vault

▼ To back up the catalog:

1. Complete the instructions in the **Attributes** area as follows:
 - a. From the **Media Server** drop-down list, select a media server.
 - b. From the **Volume Pool** drop-down list, select the volume pool from which to draw the media.

Note This volume pool should not be the same as any other volume pool being used for duplication of images.

- c. In the **Retention Period** field, enter the number of days you want to retain the catalog backup.
 2. In the Files to be backed up area you can specify exactly which catalog files or directories Vault will back up.
- NetBackup backs up certain files by default. These files will be listed in the Files area of the Catalog Backups tab. They are:

install_path\NetBackup\db - this directory contains NetBackup scheduling information, error logs, and information about files backed up from client workstations.

install_path\volmgr\database - this directory contains information about the volumes, robotics, and devices used in the current NetBackup configuration.

install_path\var - this directory contains information about licenses and authorization.

To add a file or directory to the list:

- To add a file or directory, click the **New** button.
- To delete a file or directory, click the **Delete** button.
- To change the path of a file or directory, click the **Change** button.

You must use absolute path names, for example,

UNIX: /usr/openv/netbackup/*directory_name* or *file_name*

Windows: *install_path*\NetBackup*directory_name* or *file_name*

3. The **Backup Policy...** box needs to be checked only if your catalog backup is too large to fit on one tape and you need to configure the two-stage catalog backup as documented in chapter 4 of the *NetBackup System Administrator's Guide*.
4. If you selected **Backup Policy...**, choose a backup policy name from the dropdown list box.

Note This backup policy must send the catalog data to one of your off-site volume pools which is included in the volume pool list on the Eject tab.

5. When you are satisfied with your choices, move to the next tab.

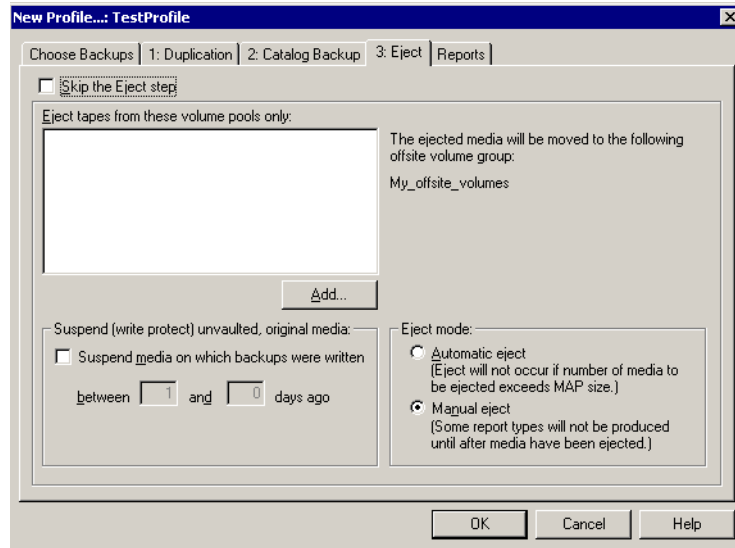
Configuring Large Catalog Backups

Large catalogs may require more than one tape for a backup. See the *NetBackup System Administrator's Guide*. The media server used for the catalog backup step must have access to the off-site volume group. You must create a new catalog backup policy for Vault. You can copy your NetBackup catalog backup policy but specify a different volume pool.



Configuring the Eject Tab

The information configured via the **Eject** tab identifies which copy of your backups are to be vaulted. For example, if you intend to vault the copies in the Vaulted-Payroll volume pool, then that volume pool would be specified in the volume pool list.



This tab enables you to:

- Edit the list of volume pools from which to eject media. Not all media in the pool are ejected -- only the media in the pool which contain images that meet the selection criteria.
- Configure Vault to automatically eject the media or to defer the eject.
- Suspend unvaulted media

▼ To set eject options:

1. Select at least one item for the list of volume pools from which tapes should be ejected.
 - a. To change the list of volume pools, click the **Add** button under the area labelled **Eject tapes from these volume pools only**.
 - b. Use the **Add/Add All** and **Remove/Remove All** buttons to move items from one category to the other.
2. Select the **Suspend media on which backups were written** checkbox if you want to suspend media on which backups were written for the specified range of days.

Since images are usually duplicated and ejected in batches, there may be backups that you would like to vault but cannot because they are not complete by the time you want to start vaulting. To prevent extraneous newer image fragments from being vaulted, suspend the tapes that contain the fragments; they will get picked up in the next batch. For more information about Suspend, see “Use the Suspend Option to Avoid Vaulting Partial Backups” on page 32.

Note Vault only suspends the tapes in volume pools specified in the Eject volume pool list.

3. Indicate whether the eject is to be automatic or manual.
 - If automatic eject is selected and a Vault session produces more media for eject than will fit in the media access port (MAP), then the Vault job will not perform an automatic eject. After the job finishes, the operator will need to manually eject the media and generate the reports. Manual eject and manual reporting must be done using `vltopmenu` or `vlteject`.
 - Manual eject requires user intervention to eject the tapes. Some reports marked with * on the Reports tab will not be generated until the media has been ejected.

If you intend to consolidate eject for multiple sessions, you should configure manual ejects. After a number of sessions have run you can use `vltopmenu` or `vlteject` to run the consolidated eject operations.
4. When you are satisfied with your choices, move to the next tab.

Configuring the Reports Tab

For information on the individual reports you can generate, see “Reporting” beginning on page 109.



▼ To complete the Reports tab:

1. If you want certain text to appear at the top of every page, type it in the **Report header** box.

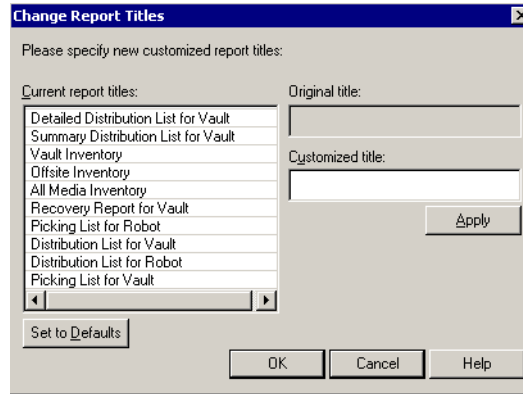
The screenshot shows the 'New Profile: MyProfile' dialog box with the 'Reports' tab selected. The 'Report header' field is empty. Under 'Reports for media going offsite', 'Picking List for Robot' is checked. Under 'Reports for media coming onsite', 'Picking List for Vault (*)' is checked. Under 'Detailed media reports', 'Vault Inventory (*)' is checked. The 'Recovery Report for Vault (*)' is unchecked, with a range of '0' to '0' days ago. The 'Report destination' section has 'Print command' checked with the value 'lpr'. The 'Report mode' section has 'Manual reports' selected. The 'Change Report Titles' button is visible.

The header is not report-specific. Therefore, you would not want to use a report title in the header area. Since report generation is a batch process, and may occur on a regular basis, including a specific date in the header is not appropriate.

2. Select the checkbox next to the reports you want to generate.

If you select **Recovery report for vault(*)**, you must set a range of days (between **N** and **n** days ago). This report shows all policies defined on a NetBackup master server and all media that is required for restores between a given set of dates. Generating this report on a regular basis will help with disaster recovery efforts.

3. Optionally, if you want to rename the report, click **Change Report Titles** to bring up the Change Report Titles dialog.



- a. Highlight the name of the report for which you want to change the name.
 - b. Specify the new name in the **Customized title** field.
 - c. Click **Apply**.
4. If you want to save the report, choose where you want to send the report in the Report destination options area, and enter the appropriate information in the field provided.
 - **Email:** Enter email addresses, separated by commas, semi-colons, or spaces. If you have already entered this information on the Email tab of the Vault Properties dialog, it will show up here automatically when you select the Email checkbox.
 - **Print Command:** Enter a print command along with a printer name. Specify the full path to the print command. If you want to add an alternate print command to print the Recovery Report in landscape format, separate the print commands with a plus (+).
 - **Folder:** Enter a path name (Windows GUI).
 - **Directory:** Enter a path name (Java GUI).
5. In the **Report mode** area, choose whether to generate the reports automatically or manually.



Note If the Report mode is set to Automatic and the Eject step is skipped, or the Eject mode is set to Manual, then the reports marked with an asterisk (*) will not be printed. You can print them using the vltopmenu MUI after the media have been manually ejected. For more information on vltopmenu, see “Using vltopmenu” beginning on page 129.

6. Click **OK**.



Inline Tape Copy lets you create up to four copies simultaneously at backup time. This chapter considers the following topics:

- ◆ Understanding Multiple Copies (Inline Tape Copy)
- ◆ Selecting Continue vs. Fail for Multiple Copies
- ◆ Using Inline Tape Copy (Multiple Copies) in Vault
- ◆ Creating Multiple Copies from Outside Vault

Understanding Multiple Copies (Inline Tape Copy)

You can create up to ten copies of a backup image. (The default is two copies.) You can create up to four copies simultaneously (if the resources allow).

Note If you are making multiple copies, all storage units must be connected to the same media server. An additional drive is required for each copy, and the destination storage units cannot be disk or optical disk. The backup time required may be somewhat longer than for one copy.

The facility to create multiple copies simultaneously (known as inline tape copy) leads to the possibility of multiple original copies. This would occur if a backup policy was configured to create multiple copies of a backup image as that backup image was first created. In addition, a vault profile can create up to four more copies of a backup image after the original backup copies have been created.

This possibility brings up several questions:

- ◆ Which copy is the primary copy? If your Vault profile does duplication, you can specify that one of the duplicate copies should be made primary. You should configure this so that the copy that is to remain on site is primary.

By default, the first successful copy is the primary copy.

- ◆ Is the expiration date of each copy the same?



The expiration date is specified separately for each copy via the Multiple Copies dialog.

- ◆ If there are multiple originals, how does Vault decide which original copy to eject?

Each of the originals should be put into a separate volume pool. The off-site volume pool for the copy to be ejected (and vaulted) must then be specified for the Eject step of the Vault profile.

Selecting Continue vs. Fail for Multiple Copies

When making multiple copies, you can choose how an operation will behave if one of the copies fails.

If You Choose Continue	If You Choose Fail
In Vault, if you choose Continue , then the copy may never get vaulted. This is because the duplicate is considered to be successful if any single copy succeeded. If all copies are marked Continue , then you are guaranteed that one copy will be successful, but you have no control over which copy is successful.	In Vault, if a copy marked Fail all copies fails, it is guaranteed that the next time the vault profile is run, Vault will again try to duplicate the image, if the following conditions are also met: <ul style="list-style-type: none">- The Vault profile's time window has enough overlap to pick up the image the next time the profile runs, and- The Vault profile does not eject and vault the original (primary) copy. If it did vault the primary copy, then the next time the profile runs, it would be unable to copy the image because the primary copy would be off site.

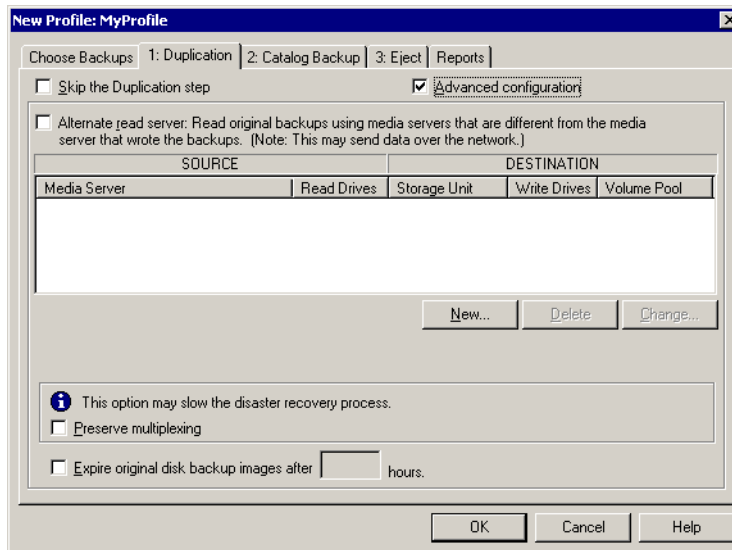


If You Choose Continue	If You Choose Fail
<p>If you choose Continue for all copies, it is likely that there will be at least one copy of the data. However, the copy you intend to vault could fail and would never be vaulted unless you took some corrective action:</p> <ul style="list-style-type: none"> Make a duplicate of the backup, ensuring that the duplicate is in the correct off-site volume pool to be ejected by the Vault profile and has the correct retention level. <p>To ensure that the image gets vaulted even if it fails during the original backup, do either of the following:</p> <ul style="list-style-type: none"> - Always watch the Activity Monitor for a failed status for the copy-to-be-vaulted. If you see a failed backup job, determine whether it's a critical copy. If so, make a duplicate of that image so that Vault can pick it up. - Configure the Vault profile to duplicate the image. This time, use the Fail all copies option. If Vault sees that there is already a copy in the destination off-site volume pool(s), it will not duplicate the image, although it would eject the appropriate copy. <p>If the copy failed during the original backup job, the Vault profile would subsequently duplicate it. If the copy succeeded during the original backup job, the vault profile would not duplicate it. Either way, the Vault profile would assure a copy, which it would then eject and send off site.</p>	<p>If you choose Fail entire job, then the entire backup job will fail and no copies will be made. In this case, normal NetBackup behavior will ensure that a successful backup for this policy eventually happens. That is, NetBackup will automatically retry the backup and/or the next time the backup window for this Policy opens up, NetBackup will again try to run the backup (regardless of the frequency of the schedule). NetBackup will do this until the backup succeeds, though it might miss one or more backup windows in the process.</p>

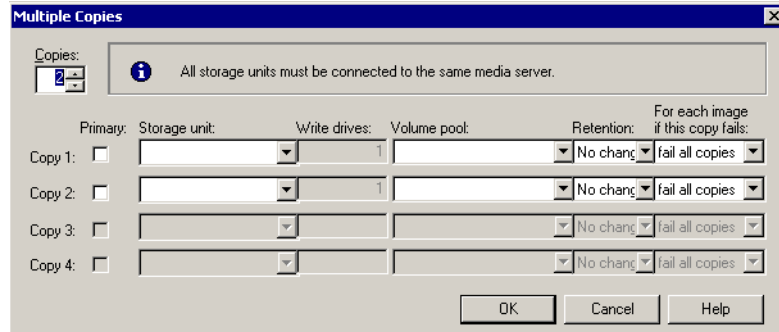


Using Inline Tape Copy (Multiple Copies) in Vault

To make multiple copies in Vault, you must select **Advanced configuration** on the Duplication tab when creating a profile:



1. If robots are shared by more than one media server, you can ask to read the original backups using a different media server from the one that wrote the backups. By default this option is turned off, since it can send data over the network and can slow performance. If you want to turn on this option, check the **Read original backups...** box. You can now specify a different media server to use for reading the original images.
2. Click **New** to bring up the Multiple Copies dialog.
3. On the **Multiple Copies** dialog, specify the storage unit to be used for the duplication. If the storage unit has more than one drive, the source and destination storage units can be the same.



4. Specify a volume pool for each copy. NetBackup does not verify in advance that the media ID selected for the duplicate copy is not the same as that of the piece of media that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure a different piece of media is used.
5. Specify the retention level for each copy. When the retention period expires, information about the expired backup will be deleted, and the backup will be unavailable for a restore.

If you select **No change**, the expiration date will be the same for the duplicate and original copies. You can then use `bpexpdate` to change the expiration date of the duplicate copy if you wish.

If you select a different retention period, the expiration date of the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2001, and its retention period is one week, the new copy's expiration date is November 21, 2001.

6. Indicate what action is to be taken if a copy fails. By default, all subsequent copies will fail. You can also choose to continue the duplication process if one copy fails.
 - If you choose **Fail all copies**, the next time this profile is run Vault will try to duplicate the image again, if the following conditions are met:
 - The vault profile's time window has enough overlap to pick up the image the next time the profile runs, and
 - The vault profile does not eject or place the original (primary) copy of the backup into an off-site volume group. If this is the case, then the next time the profile runs it will not be able to copy the backup because the backup will be inaccessible.



- If you choose **Continue**, a copy of the backup may never be placed in the off-site volume group because Vault will consider the duplication job successful if any of the copies succeed. However, the successful copy may not be the copy designated for off-site vaulting.
7. Indicate which copy is to be primary, if you want one of these four copies to be the primary copy. The primary copy is the copy that is used for restores. The original copy is the primary copy by default. If the copy that you indicate as primary fails, and you have set Multiple Copies to continue in case of failure, the first successful copy is the default primary copy. See “What Is the Primary Copy?” on page 5 for more information.
 8. Click **OK**.
 9. Indicate whether you want to preserve multiplexing.

Multiplexing is the process of sending concurrent-multiple backup images from one or more clients to the same piece of media. This process will speed up duplication, but will slow down restores and disaster recovery processes. The multiplexed duplication process will occur for all multiplexed images that are selected for duplication during a given Vault session. For more information about multiplexing, refer to the *NetBackup System Administrator's Guide*.

Note If multiplexing is enabled in Vault, make sure that all destination storage units have multiplexing enabled. (Set through Media Manager -- **Global Configuration of Storage Unit**.)

10. Check **Expire original...** if you want to set an expiration time, and select the number of hours.

Each tape has an expiration date assigned to it. You can use Expire Original to force an earlier expiration time so the tape can be reused. Be sure you allot enough time for the duplication to be completed.
11. When you have filled in the appropriate information, move to the next tab.

Creating Multiple Copies from Outside Vault

Outside Vault, you can configure Inline Tape Copy to make multiple copies in the following ways:

- ◆ Through Policy node
- ◆ Through the Catalog node

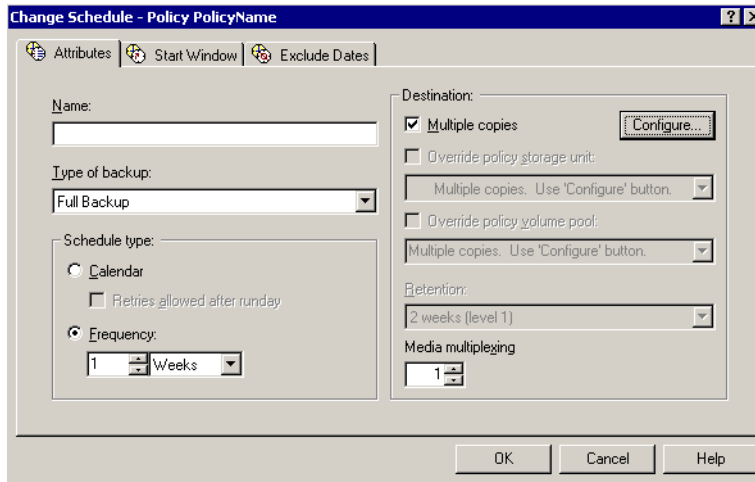
Configuring Inline Tape Copy through the Policy Node

Use the following procedure to configure Inline Tape Copy through the Policy node.

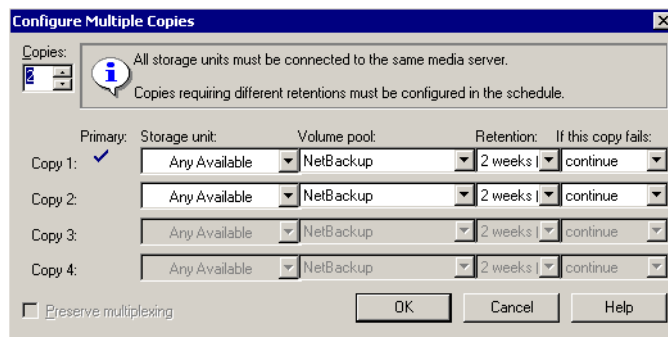
▼ To configure multiple copies for Inline Tape Copy

1. In the NetBackup Administration Console, expand **Master Server > NetBackup Management > Policies**.
2. Double-click an existing policy or:
 - Windows: select **Actions > New > New Policy**
 - UNIX: click the **Add new policy** button to create a new policy.
3. Select the Schedules tab to configure Inline Tape Copy.
4. Double-click an existing schedule or click **New** to create a new schedule.
5. In the Schedule Attributes tab, select **Multiple copies**, then click **Configure**.





6. Specify the number of copies to be created simultaneously. (The maximum is four.) Copy 1 is the primary copy. If Copy 1 fails for some reason, the first successful copy is the primary copy.



7. Specify the storage unit where each copy will be stored. (This is not applicable to disk type storage units.) If a storage unit has multiple drives, it can be used for both the source and the destination.
8. Specify the volume pool where each copy will be stored.
9. Select the retention level for each copy. For more information on setting retention levels, see the *NetBackup System Administrator's Guide*.

10. In the event that the copy should not complete, select whether you'd like the entire job to fail, or whether you'd like the other copy (or copies) to continue.
11. Click OK.

Configuring Inline Tape Copy Through Catalog Node

NetBackup can create up to 10 copies of unexpired backups.

If licensed to do so, NetBackup can create up to four of the copies simultaneously. (See the note below.) The number of backup copies is determined by the Host properties Global Attributes setting, **Maximum Backup Copies**. See the *NetBackup System Administrator's Guide* for more information.

Duplicating requires a minimum of two drives:

- ◆ one drive to read the original
- ◆ one drive to create the copy

Note An alternative to duplicating backups is to use Inline Tape Copy. Inline Tape Copy allows you to create up to four copies simultaneously at backup time. Keep in mind that an additional drive is required for each copy and the destination storage units cannot be disk or optical disk. The backup time may be somewhat longer than for one copy only.

NetBackup does not verify in advance whether the storage units and drives required for the duplicate operation are available for use, only that the destination storage unit exists.



Where Duplication Is and Is Not Possible

The following lists describe scenarios which present candidates for duplication and scenarios where duplication is not possible:

Possible to duplicate backups:	Not possible to duplicate backups:
<ul style="list-style-type: none">- from one storage unit to another.- from one media density to another.- from one server to another.- from multiplex to nonmultiplex format.- from multiplex format and retain the multiplex format on the duplicate. The duplicate can contain all or any subset of the backups that were included in the original multiplexed group. This is done with a single pass of the tape. (A multiplexed group is a set of backups that were multiplexed together during a single session.)	<ul style="list-style-type: none">- while the backup is being created (unless when using Inline Tape Copy).- when the backup has expired.- using the NetBackup scheduler to automatically schedule duplications.- of the NetBackup catalogs.- when it is a multiplexed duplicate of the following:<ul style="list-style-type: none">– Auspex FastBackup– FlashBackup– NDMP backup– Backups from disk type storage units– Backups to disk type storage units– Nonmultiplexed backups

Note Do not duplicate images while a NetBackup catalog backup is running. This results in the catalog backup not having information about the duplication.

Notes on Multiplexed Duplication

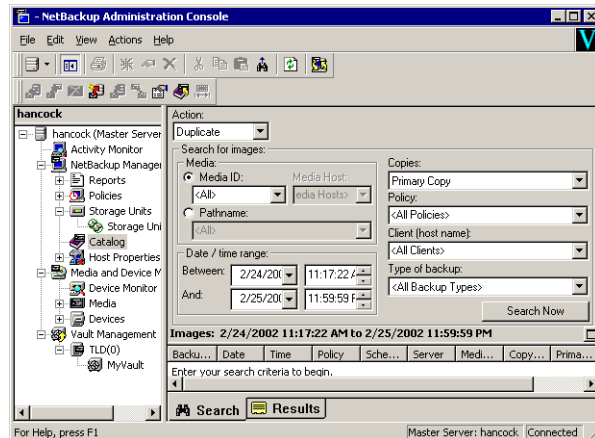
- ◆ When duplicating multiplexed SQL-BackTrack backups with multiplex mode enabled, it is necessary to duplicate all the backups in the multiplexed group. This ensures that the fragment order and size is maintained in the duplicate. Otherwise, it is possible that restores from the duplicated backups will not work. A multiplexed group is a set of backups that were multiplexed together during a single multiplexing session.
- ◆ When duplicating multiplexed backups, the multiplex settings of the destination storage unit and the original schedule are ignored. However, if multiple multiplexed groups are duplicated, the grouping within each multiplexed group is maintained. This means that the duplicated groups will have a multiplexing factor that is no greater than used during the original backup.
- ◆ If all backups in a multiplexed group are duplicated to a storage unit that has the same characteristics as the one where the backup was originally performed, the duplicated group will be identical, with the following exceptions:

- If EOM (end of media) is encountered on either the source or destination media.
- If any of the fragments in the source backups are zero length (this can occur if many multiplexed backups start at the same time), then during duplication these zero-length fragments are removed.

Note This is important only for SQL-BackTrack backups.

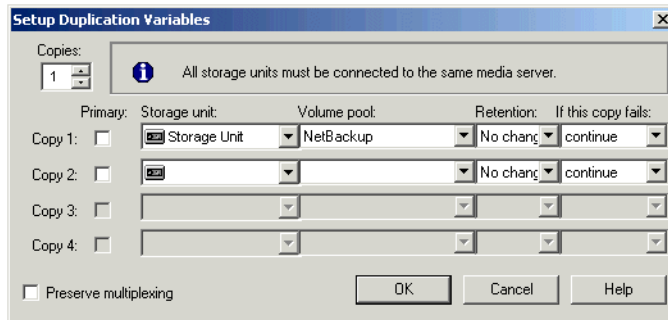
▼ **To duplicate backup images through the Catalog node:**

1. In the NetBackup Administration Console, expand **Master Server > NetBackup Management > Catalog**.
2. Set up the search criteria for the image you wish to duplicate. Click **Search Now**.
3. Right-click the image you wish to duplicate and select **Duplicate** from the shortcut menu. The **Setup Duplication Variables** dialog appears.



4. Specify the number of copies to be created.





If the Inline Tape Copy option or NetBackup Vault is installed and there are enough drives available, the copies will be created simultaneously. Otherwise, the system may require operator intervention if, for instance, four copies are to be created and there are only two drives.

5. The primary copy is the copy from which restores will be done. Normally, the original backup will be the primary copy.

If you want one of the duplicated copies to become the primary copy, then check the appropriate box. Otherwise leave the fields blank.

6. Specify the storage unit where each copy will be stored. If a storage unit has multiple drives, it can be used for both the source and destination. Disk type storage units may be used if only one copy is to be made.

7. Specify the volume pool where each copy will be stored.

NetBackup does not verify in advance that the media ID selected for the duplicate copy is not the same as the media ID of the volume that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure a different volume is used.

8. Select the retention level for the copy, or leave it the same.

The duplicate copy shares many attributes of the primary copy, including backup ID. Other attributes, such as elapsed time, apply only to the primary. It is the primary copy that NetBackup uses to satisfy restore requests.

- If **No change** is selected for the retention period, the expiration date is the same for the duplicate and source copies. You can use the `bpexpdate` command to change the expiration date of the duplicate.

- If a retention period is indicated, the expiration date for the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2001 and its retention period is one week, the new copy's expiration date is November 21, 2001.
9. Specify whether the remaining copies should continue or fail if the specified copy fails.
 10. If the selection includes multiplexed backups and the backups are to remain multiplexed in the duplicate, check **Preserve multiplexing**. If you do not duplicate all the backups in a multiplexed group, the duplicate will have a different layout of fragments. (A multiplexed group is a set of backups that were multiplexed together during a single session.)

By default, duplication is done serially and attempts to minimize media mounts and positioning time. Only one backup is processed at a time. If **Preserved multiplexing** is enabled, NetBackup first duplicates all backups that cannot be multiplex duplicated before the multiplexed backups are duplicated.
 11. Click **OK** to start duplicating.
 12. Click the **Results** tab, then select the duplication job just created to view the job results. For more information, see the *NetBackup System Administrator's Guide*.





This chapter covers administrative tasks in Vault and situations where you might want to make a change to your Vault configuration.

This chapter covers the following topics:

- ◆ What Is a Vault Session?
- ◆ Entering Addresses for Email Notification
- ◆ Editing a Vault or Profile
- ◆ Copying a Profile
- ◆ Moving a Vault to a Different Robot
- ◆ Adding Alternate Media Server Names
- ◆ Understanding Log Files
- ◆ Using Notify Scripts
- ◆ Vault Support in Activity Monitor
- ◆ Ensuring Available Media for Catalog Backups

What Is a Vault Session?

A Vault session, or vaulting job, is the process of executing the steps specified in a Vault profile. A profile may consist of any or all of these steps: duplicating images, backing up the catalog, and ejecting media. A profile also indicates which reports should be generated at the end of the Vault session.

Vault can have several vault jobs running on the same server simultaneously as long as they have different vault names and use different volume pools. This assumes your master server has the resources to do this and the configuration allows it.

Prior to running a vault session, at least one robot, one vault, and one profile must have been configured.



Running a Vault Session

When you start a Vault session, you begin to perform the tasks you have configured in the selected profile. A vault session can be run in any of these ways:

- ◆ From the Vault GUI
- ◆ From the Vault Administration Menu User Interface (`vltadm`)
- ◆ By a NetBackup policy
- ◆ Directly from the command line

The section below, “Running a Vault Session from the GUI”, describes how to start and run a session from the GUI and from the NetBackup Scheduler.

Only one instance of the GUI or MUI should be run at one time because they both refer to the same files for configuration information. One instance of any of these applications can overwrite the output of another.

To run a vault session from the command line, set your environment variable `PATH` to contain the path in which the NetBackup binaries are installed. You must specify the robot number, vault, and profile; or you can specify just the profile, if it has a unique name. For example:

```
vltrun 1/your_vault/your_profile
```

where *1* is the robot number, *your_vault* is the vault, and *your_profile* is the profile.

Note In NetBackup Vault 4.5, running a vault session from a MUI is only possible on UNIX systems. Refer to “Using `vltadm`” on page 118 for more information.

Previewing a Vault Session

The preview option lets you sanity check the configuration for a profile. After you have run the preview option, check the results in the `preview.list` file. This file contains information about each image selected through the criteria you have specified. The list of images generated by preview may be a superset of the backups that will actually be vaulted if you do not further narrow the selection. That is, under certain circumstances, the list will contain more backup images than will be vaulted:

- ◆ If the profile is configured to duplicate only disk images, then selected images on removable media will not be vaulted.
- ◆ If there are any backups in the list that do not have a copy on media in one of the volume pools listed for the eject step, they will not be vaulted.

You can preview a Vault session from the command line, using the `vltrun` command with the `-preview` option. You must specify either a unique profile, or the `robot_number/vault_name/profile` parameter.

For example:

```
vltrun 1/your_vault/your_profile -preview
```

where *1* is your robot, *your_vault* is your vault, and *your_profile* is your profile.

The `-preview` option starts a new vault job, performs a search on the image catalog based on the criteria specified on the Choose Backups tab, and then exits. Vault does not act on the images selected. To view the list of selected images, display the `preview.list` file, located in:

UNIX: `install_path/netbackup/vault/sessions/vault_name/sidxxx`

Windows: `install_path\NetBackup\Vault\sessions\vault_name\sidxxx`

The list of images generated by the `-preview` option is a superset of the images that will actually be vaulted. Two more criteria may be applied to the list before the actual list of images to be vaulted is finalized. For example, if only disk images are to be duplicated, then any backups which appear in the list but do not have a copy on removable media will not be vaulted. Or, if the Eject step does not list the off-site volume pool for some of the selected images, they will not be vaulted.

Running a Vault Session from the GUI

▼ To run a Vault session from the GUI:

1. Start the NetBackup Administrative GUI.

2. Expand **Vault Management** in the left pane.

The names of the robots configured for NetBackup display.

3. Pick the robot you want to use for this session.

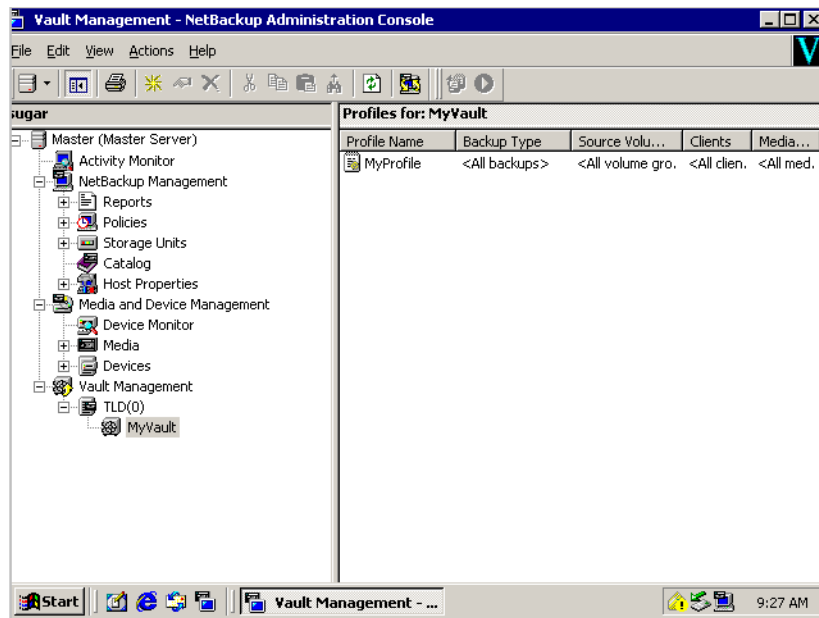
The names of the vaults configured for the robot display.

If no vaults display, right-click on the robot and choose **New vault** to configure a vault. See “Creating a Vault” on page 47 for more information.

4. Highlight the vault you want to use for this session.

The profiles for this vault display in the right pane.





If no profiles display, right-click on the vault name in the left pane and choose **New profile** to configure a profile. See “Creating a Profile” on page 50 for more information.

5. Highlight the profile you want to use.
6. Right-click on the profile to bring up a context menu with the available actions for this profile.

One of the options is to rearrange the columns displayed for the profile. Two columns of information are not displayed by default; these are **Volume Pools (Ejected)** and **Report Destination**. If you wish to display these columns, click **Columns...**, then **Layout**, and rearrange the current column display.

7. Select **Start Session**.

Start Session remains highlighted until the session begins. Once the session starts, the GUI displays the message:

Manual vault session for profile has been started. Use the Activity Monitor to view progress

Running a Session using the NetBackup Scheduler

To run a vault session using the NetBackup Scheduler, configure a Vault policy. You can also run the vault session manually in this case, by selecting the option to run a manual backup after you have configured the Vault policy. Refer to the *NetBackup System Administrator's Guide* for information on how to configure and run a policy.

Configuration details for a Vault policy:

- ◆ Select Vault as the policy type.
- ◆ Select Automatic Vault as the schedule type.
- ◆ Add the following to the include list: a command line for the vltrun command, specifying the robot, vault, and profile for the session. The include list will look something like this:

UNIX: `/usr/opensv/netbackup/bin/vltrun
robot_number/vault/profile`

Windows: `install_path\NetBackup\bin\vltrun
robot_number/vault/profile`

Enclose the `robot_name/vault_name/profile` parameter in double quotes if it contains unusual characters. If the profile name is unique, the robot number and vault name may be omitted.

Resuming a Vault Session

If your vault job fails or returns a Partially Successful completion status, begin by evaluating the cause of the problem. If the Activity Monitor or the emailed notification of session status (`netbackup/vault/sessions/vault/sidxxx/summary.log`) does not contain enough information, examine the log files (see the Troubleshooting chapter).

Then address the cause of the problem.

- ◆ If the session had reached the Eject step or had attempted to generate reports before encountering problems, then you can simply use `vltopmenu` (or `vlteject`) to finish the eject and/or reporting for the session.
- ◆ Otherwise start a new session for your profile. If you are doing duplication, Vault will not re-duplicate any images it already duplicated, but it will eject those images.



Entering Addresses for Email Notification

An email notification is sent at the end of the vault session to the email address or addresses configured in the Vault Properties dialog. The mail message provides a summary of the vault session, in the form of a `summary.log` file, along with the completion status. The subject of the mail message is formatted as follows:

Vault on *Master Server*: Status Code [*robot_number/vault_name/profile*]

Where:

Master Server = hostname of the master server

Code = Error code

Robot = Robot number

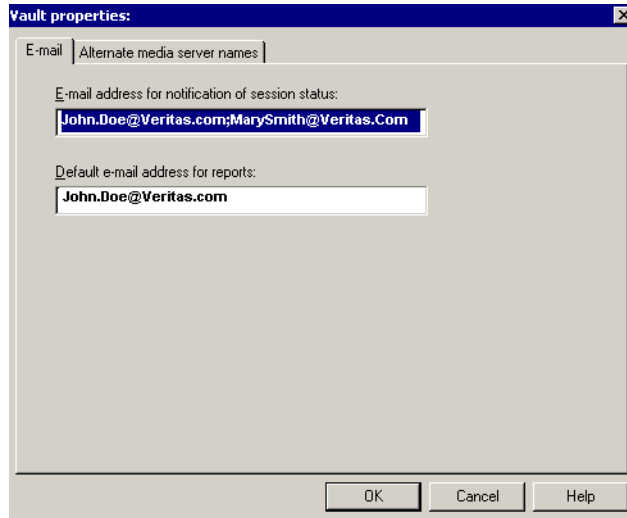
Vault = Vault name

Profile = Profile name

Note Prior to setting up email notification in Vault: On UNIX systems the mail or mailx command is used. On Windows systems, the email client must be configured through the `nbmail.cmd` script.

▼ To set up email notification:

1. From within Vault Management, open the **Actions** menu and select **Vault Properties**. The Vault Properties dialog will display.
2. Select the **Email** tab.



3. Enter an email address in the **...for notification of session status** field if you want certain people, such as the Vault administrator, notified of success or failure when a Vault session completes.
4. Enter an email address in the **...email address for reports** field if you want certain people to receive Vault reports through email.

If you want to enter more than one address in either field, separate the addresses with commas.

Editing a Vault or Profile

If you want to edit a vault or a profile, right-click on the item and select **Change**. Then complete the dialog and click **OK**.

Printing Vault and Profile Information

You can print a list of the information (robots, vaults, or profiles) that is currently displayed in the right pane for reference. From the **File** menu, choose **Print**, or click the **Print** icon on the toolbar.



Copying a Profile

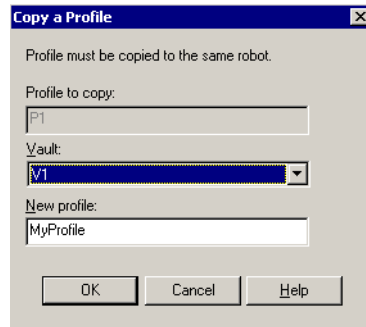
If you want to create a profile that is similar to another profile, you can save effort by copying the profile, renaming it, and making whatever changes are necessary.

Note The new profile must belong to the same robot as the original profile.

▼ **To copy a profile:**

1. Highlight the profile you want to copy.
2. Open the **Actions** menu and select **Copy Profile**.

The Copy a Profile dialog displays.



3. From the dropdown list under the **Vault** field, choose the vault you want the new profile to be part of.
4. Type a new name for the profile.
5. Click **OK**.

Moving a Vault to a Different Robot

A vault belongs to a particular robot.

Robots are configured under NetBackup. The action described here does not change the configuration of a robot in Media Manager. You cannot edit a robot. However, you can change which robot the vault is part of. To do so, right-click on the robot and select **Change**. Then complete the dialog, specifying another robot for the vault, and click **OK**.

Note All vaults that had been associated with the old robot will now be associated with the one chosen in the dialog. Some profile configurations may be invalid under the new robot, for example, if the old robot was associated with a media server that the new robot is not.

Adding Alternate Media Server Names

If you change the name of a media server, you may want to create a list of server names to ensure that the profile will capture images for this media server under all of its names.

Note If you add more than one name, the names for a single server must be on the same line, separated with commas.

If you specify names that belong to more than one server on the same line, Vault will assume that the names belong to the same server, and will act accordingly. This use of the Alternate Media Server Names dialog allows you to specify the same destination storage unit for two (or more) different servers. This is useful if you want to duplicate images from a number of servers.

Caveats:

- ◆ You must know that you have enough drives in the specified destination storage unit to keep up with the demand for duplication. If you do not, you risk a deadlock situation.
- ◆ You must know that the specified media servers have access to the destination storage unit. If not, again, you risk a deadlock situation, and your Vault job will fail. To prevent this situation, use the **Media Servers** criterion on the **Choose Backups** tab to ensure that only backup from certain media servers will be selected.
- ◆ You need to take into account the fact that depending on the media server(s) in use, this configuration can send data over the network.

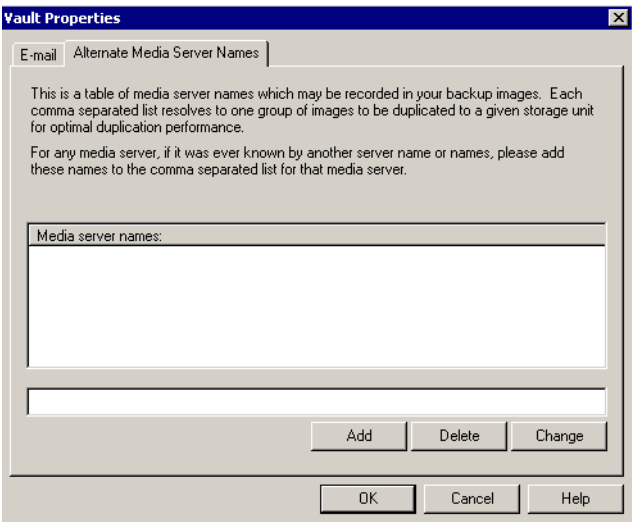
Note We recommend that you only specify one destination storage unit per server. If you specify more than one, you may create a problem as Vault does not have a mechanism to choose which destination storage unit to send the duplicate images to.

▼ To add alternate media server names:

1. From anywhere within Vault, open the **Actions** menu and select **Vault Properties**.
By default, the **Email** tab of the Vault Properties displays.



2. Select the **Alternate media server names** tab.



By default, the **Media Server Names** area will be blank, unless you have already added one or more names.

3. In the field below the **Media Server Names** area, type the alternate name for the media server, and click **Add**.
- If you want to remove a name you previously added, highlight it and click **Delete**.
 - If you want to change a name you previously added, highlight it and click **Change**.

Understanding Log Files

Vault generates two types of logs: session logs and debug logs. These files help you keep track of Vault processes and troubleshoot when necessary.

The following table summarizes the log options.

Name of Log File	Location of Log File	Purpose of Log File
duplicate.log. <i>nn</i>	../netbackup/vault/ sessions/vault_name/sidxxx	Generated by -L option of the bpduplicate command.



Name of Log File	Location of Log File	Purpose of Log File
<code>preview.list</code>	<code>../netbackup/vault/ sessions/vault_name/sidxxx</code>	Summary of images to be duplicated if Duplication step is configured, or ejected, if Eject step is configured and Duplication step is not.
<code>image.list</code>	<code>../netbackup/vault/ sessions/vault_name/sidxxx</code>	Lists all images and partial images for a session.
<code>detail.log</code>	<code>../netbackup/vault/ sessions/vault_name/sidxxx</code>	Summary of each action performed for a Vault session.
<code>summary.log</code>	<code>../netbackup/vault/ sessions/vault_name/sidxxx</code>	Brief description of the Vault session and its results. If email notification is enabled, data in this log file is appended.
<code>vltrun. output_vault _name</code>	<code>../netbackup/vault/ sessions/vault_name/ sidxxx</code>	Output file produced by Vault process showing progress of Vault session.
<code>log.mmddyy</code>	<code>../netbackup/logs/vault</code>	Debug log for Vault sessions run on a particular day.
<code>duplicate.err .nn</code>	<code>../netbackup/logs/vault</code>	Contains the stderr and stdout of the <code>bpduplicate</code> command. Usually empty unless there is a serious failure in <code>bpduplicate</code> .

◆ Session Logs

The session directory generated for each vault session collects information for the session in two log files. The file `detail.log` contains a step-by-step account of each action performed for the session. Some of the information in `detail.log` is also recorded in the NetBackup log files. The `summary.log` file contains a brief description of the vault session, and the results of the session. If email notification is enabled, the information in this file is appended to the email.

The `detail.log` has information about the number of images selected by a particular session. In addition, it should record information (during the duplication step) about the total number of images and the number of images duplicated. If these numbers do not match, it means that some images were not duplicated. The log should contain information about which images were not duplicated, either because



they were duplicated in a prior session or because the duplication failed for some reason. The actual images selected by the session will show up only if a higher debug level (level 5) is used.

Vault maintains its session log files for a particular session in the directory for that session. The directory is located in the following path:

- UNIX:
`/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx`
- Windows:
`install_path\NetBackup\vault\sessions\vault_name\sidxxx`
- where *vault* is the name of the vault used for this session and

where *xxx* is the unique session identifier that Vault assigns to each vault session. This value starts at 1 the first time Vault runs, and is incremented by one for each new session. The session identifier for a particular vault can be found by looking at the Activity Monitor entry for that session.

◆ Debug Logs

Debug logging for Vault uses the standard NetBackup debug logging path. Vault debug logs are stored in the following path:

- UNIX: `/usr/opensv/netbackup/logs/vault`
- Windows: `install_path\NetBackup\logs\vault`

Within this directory, each day's logs are stored in `summary.log` and `detail.log`.

You must create a subdirectory for the specific NetBackup software you are interested in monitoring. For example, you can monitor the process `bptm` to see the tape manager output. Create the directory on the server where the process runs.

Other debug logs that may be useful in tracking a vault session are the `bpsched` log and the `bprmvlt` log. The `bprmvlt` process submits and tracks a scheduled vault job. It is the Vault equivalent of `bpbrm` for a backup job.

For more information about these log files, see the *NetBackup System Administration Guide* or the *NetBackup Troubleshooting Guide*.

Output from the Vault Driver

The vault driver, `vltrun`, produces an output file showing the progress of the vault session. This file resides in the each vault session directory. It is called `vltrun.output_vault_name`.

As each step of a vault session completes, the result of the step is written to this file. If the session fails, the file only contains information up through the last step successfully completed.

Setting the Duration of Vault Working Files

Vault's working files are stored in the following directory:

UNIX: `/usr/opensv/netbackup/vault/sessions`

Windows: `install_path\NetBackup\vault\sessions`

Set the length of time NetBackup keeps these files in the NetBackup Administration console or in `bp.conf`.

▼ To set the duration of time to keep working files:

1. In the NetBackup Administration Console, select **Host Properties**.
2. Select **Master Server** under **Host Properties**.
3. In the right pane, right-click the master server and choose **Properties (Read/Write)**.
4. Select **Global NetBackup Attributes**.
5. Set the length of time to retain the Vault working files:
 - In the UNIX GUI: In the **Keep Vault Logs For** field, set a length of time to retain Vault working files. The default is 30 days. A value of 0 means the files will be kept forever.

This corresponds to the `bp.conf` entry:

```
KEEP_VAULT_SESSIONS_DAYS = days
```

 - In the Windows GUI: In the **Days to Keep Vault Files** field, set a length of time to retain Vault working files. The default is 30 days. A value of 0 means the files will be kept forever.

When the set time has elapsed, the entire `sidxxx` directory is deleted.

You should plan to retain each `sidxxx` directory at least as long as the period of time over which you plan to span consolidated ejects. We suggest that you keep these directories at least a week longer than the consolidation span. If the `sidxxx` directory has been deleted, Vault will be unable to eject tapes or generate reports from that session.

Setting the Duration and Level of Vault Log Files

Vault's log files are stored in the following directory:

UNIX: `/usr/opensv/netbackup/logs`

Windows: `install_path\NetBackup\logs`



Use the NetBackup Administration console or `bp.conf` to set the length of time NetBackup keeps these logs and the level of information contained in them.

▼ **To set the duration and level for log files:**

1. In the NetBackup Administration Console, select **Host Properties**.
2. Select **Master Server** under **Host Properties**.
3. In the right pane, right-click the master server and choose **Properties (Read/Write)**.
4. Select **Global NetBackup Attributes**.
5. In the UNIX GUI: In the **Keep Logs For** field, set a length of time to retain Vault logs. This setting applies to all NetBackup logs, including but not limited to, the Vault logs.

In the Windows GUI: In the **Duration to Retain Logs** field, set a length of time to retain NetBackup logs. This setting applies to all NetBackup logs, including but not limited to, the Vault logs.

6. Select **Logging** to set the logging level.

The logging level determines how much information is displayed in the log. A level of 0 will display the minimum amount of information; a level of 5 will display the maximum.

- In the UNIX GUI: To change the logging level, select the **Vault Logging Level** field and use the arrows to choose a different logging level. To get the most information, use level 5.

For UNIX systems, the Logging level field corresponds to the `bp.conf` entry:

`VAULT_VERBOSE = level`

- In the Windows GUI: To change the logging level, select **Vault Logging** and choose the level from the dropdown list box. To get the most information, use level 5.

Using Notify Scripts

The vault job can call shell scripts at well-defined checkpoints. You can customize these scripts to perform site-specific processing. These scripts are installed in the following directory:

UNIX: `/usr/openv/netbackup/bin/goodies`

Windows: `install path\NetBackup\bin\goodies`

They must be copied to the `bin` directory for them to be executed.

The scripts must return a normal status (0) for the vault job to continue processing. In case of failure, the script must return a nonzero status code to cause the vault job to stop. On UNIX systems, the return status is communicated to the vault job through the `exit` call. On Windows systems, the scripts communicate the return status in a file. The path name of the file is communicated by the vault job to the script through the environment variable `RESULT_FILE`.

The following scripts have been provided with Vault:

◆ `vlt_start_notify`

This script is called by the vault session after it starts. The input options that `vltrun` provides when executing this script are the robot name, vault name, profile name and the session id.

◆ `vlt_end_notify`

This script is called by the vault session just before it exits. The input options that `vltrun` provides when executing this script are the robot name, vault name, profile name, session id and the completion status of the vault job.

◆ `vlt_ejectlist_notify`

This script is called by the vault session just before vault tapes are ejected. This script is intended to provide a means to use vault to eject non-NetBackup tapes that are managed by Media Manager. Non-NetBackup tapes can be vaulted with NetBackup tapes or without NetBackup tapes.

The script must be modified to add the list of non-NetBackup media ids to the `addon_medialist` file. The media ids in this file will be added to the list of NetBackup media to be ejected by the current vault session. Vault will assign vendor slot ids to these tapes and track them along with the NetBackup tapes.

The input options that `vltrun` provides when executing this script are the robot name, vault name, profile name and the session id.

◆ `vlt_endeject_notify`

This script is called at the end of eject processing. The input options that `vltrun` provides when executing this script are the robot name, vault name, profile name and the session id.

These template scripts (`vlt_start_notify`, `vlt_end_notify`,...) can be customized for your entire vault configuration, in which case Vault will execute the same script for every *robot_number/vault_name/profile* combination.

Vault also lets you customize each notify script for each profile or vault or robot in your configuration.

The notify scripts can be created for:



- ◆ a specific *robot_number/vault_name/profile* combination
- ◆ a specific *robot_number/vault_name* combination
- ◆ a specific robot

Notify Script for a Specific Profile

For example, to create a `vlt_start_notify` script for a specific robot/vault/profile combination, copy `/usr/opensv/netbackup/bin/goodies/vlt_start_notify` to:

UNIX: `/usr/opensv/netbackup/bin/vlt_start_notify.robot_number.vault_name.profile_name`

Windows: `install_path\NetBackup\bin\vlt_start_notify.robot_number.vault_name.profile_name`

This script will be executed for a specific profile defined for a specific vault. This allows you to create a unique customized script for every profile in your configuration.

Notify Script for a Specific Vault

You can also create a script for a specific vault. For example:

UNIX: `/usr/opensv/netbackup/bin/vlt_start_notify.robot_number.vault_name`

Windows:
`install_path\NetBackup\bin\vlt_start_notify.robot_number.vault_name`

This script will be executed for all profiles created for a specific vault. This allows you to have a unique customized script for each vault in your configuration.

Notify Script for a Specific Robot

To have a common script for a specific robot, you need to create:

UNIX: `/usr/opensv/netbackup/bin/vlt_start_notify.robot_number`

Windows:
`install_path\NetBackup\bin\vlt_start_notify.robot_number`

This script will be executed for all profiles created for the robot. This allows you to create a unique customized script for each robot in your configuration.

Order of Execution

The order in which these scripts are executed is as follows:

1. If `vlt_start_notify.robot_number.vault_name.profile_name` is present, it will be executed, otherwise go to step 2.
2. If `vlt_start_notify.robot_number.vault_name` is present, it will be executed, otherwise go to step 3.
3. If `vlt_start_notify.robot_number` is present, it will be executed, otherwise go to the step 4.
4. If `vlt_start_notify` is present, it will be executed, otherwise go to step 5.
5. No script will be executed when the session starts. Vault will continue processing normally.

Note The customization of scripts per robot, vault and profile applies to all notify scripts documented on both UNIX and NT.

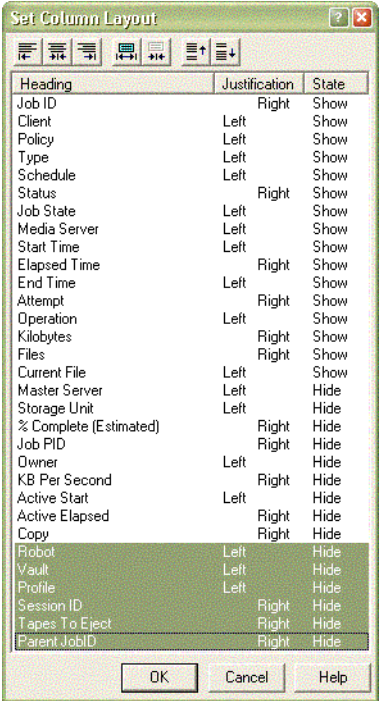
Vault Support in Activity Monitor

The Activity Monitor displays vault jobs as they execute and after they complete their execution. For a vault job initiated by the NetBackup scheduler using a policy with policy-type Vault, the policy name appears in the **Policy** field of the Activity Monitor. If the vault job is initiated from the command line or by the Vault GUI or any means other than the NetBackup scheduler, the **Policy** field is empty.

The following fields display Vault job attributes in the Activity Monitor GUI.

Note The Activity Monitor will not display these Vault fields by default. Refer to the Activity Monitor documentation to select the correct set of fields for your configuration.





:

Field	Description
Robot	The name of the robot the vault is associated with.
Vault	The name of the vault under which this session is running.
Profile	The name of the profile that holds the processing information for a vault session.
Session ID	This is the session ID, a unique numeric value, for this vault session. Session ID assignment starts at 1 the first time a vault session is run after vault has been installed. The value is incremented by one every time a new vault session runs. If a vault job is restarted, the restarted session uses the session ID of the original job.
Tapes to Eject	The number of tapes to be ejected for a vault session. Note If the Vault profile is configured for manual eject, the tapes may not have been ejected yet.

Field	Description
Parent JobID	A Vault job starts one or more bpduplicate processes. Each of these child jobs refers to the job ID of the Vault job (parent) that started it.
Operation	For vault jobs, the field contains one of the following values. These values progress from the first value to the last as the vault job progresses: Choosing Images Duplicating Images Choosing Media Catalog Backup Eject and Report

Stopping a Vault Session

You can use Activity Monitor to stop a Vault session.

▼ To stop a vault session:

1. In the Activity Monitor, highlight the vault session you want to stop.
2. From the **Action** menu, select **Cancel Job**.

Note If a vault session fails, you cannot run a new session until the old session has ended. Use **Cancel Job** to end the failed session.

Extended Error Codes

Vault jobs may exit with exit-status values greater than 255. These values are called extended error codes because they extend beyond the standard 255 NetBackup error codes. If a vault job exits with an extended error code, the exit status returned to the shell is 252. NetBackup has adopted the convention that the exit status 252 means that an extended error code is returned via stderr, in this message:

```
EXIT status = extended error code
```

The Activity Monitor displays the extended error code, rather than the value 252 returned to the shell, in this case. For more information about error codes in Vault, please see “Errors Returned by the Vault Session” on page 135 in the Troubleshooting chapter.



Ensuring Available Media for Catalog Backups

The Catalog Backup operation is different from a typical backup operation in a number of ways in how it handles media:

- ◆ If there are no unassigned media available in the catalog backup volume pool, it will not retrieve unassigned media from the scratch pool for doing a Catalog Backup (as it would have done for a typical backup operation). Therefore, if the profile is configured to do a catalog backup, make sure that there is either an unassigned tape or an old, expired catalog backup tape in the catalog backup volume pool before beginning the Vault session.

Otherwise no tape will be available and the Catalog Backup step will fail. In this case, the Vault session will continue to run and will likely complete with a Partially Successful completion status. If you then add an unassigned tape to the catalog backup volume pool and start a new session for the profile, then the new session may do nothing more than run the catalog backup, ejecting only the catalog backup tape.)

- ◆ Catalog backup tapes are also not expired in the same way that regular backup tapes are. When Vault does a catalog backup, it will use the retention period configured for the Catalog Backup step to calculate a timestamp for expiring the Catalog Backup tape. That timestamp is written to the return date MM field for the tape. When that time has passed, Vault considers the catalog backup tape to be expired and ready for reuse, in which case Vault will automatically deassign and reuse the tape.

Note If you have a very large catalog and are using the multi-tape catalog backup method, the volume pool specified in the backup policy that backs up the image catalog must not be the same volume pool specified for the regular Catalog Backup operation. Rather, it must use a volume pool included in the volume pool list for the Eject step. These tapes are handled just like any other tape used by a backup policy and should use a volume pool used by other backup policies.

Manual Deassigning of Vaulted NetBackup Catalog Media

To remove the NetBackup catalog tapes from Vault's control, you must deassign each tape. You would do this if you wanted to reuse a tape and it hadn't expired yet. After you execute the following command, you can use the tape again for a catalog backup. If you want to use the tape as a general NetBackup backup tape, you must also relabel the tape using `bplabel`. Refer the NetBackup *System Administrator's Guide* for information on `bplabel`.

The following command should be used with caution:

```
UNIX: /usr/opensv/volmgr/bin/vmquery -deassignbyid media-id
pool-num 1
```

```
Windows: install_path\volmgr\bin\vmquery.exe -deassignbyid
media-id pool-num 1
```

In the above command, *media-id* is the media ID that is to be deassigned, and *pool-num* is the pool number that *media-id* belongs to. To obtain the pool number, use the `vmquery` command as in the example below:

```
(UNIX)# /usr/opensv/volmgr/bin/vmquery -m S04440
(Windows) #install_path\volmgr\bin\vmquery.exe -m S04440
=====
media ID:S04440
media type:8MM cartridge tape (4)
barcode:-----
description:CH_V1|101|S278|00000000
volume pool:Offsite_Catalog (3)
robot type:NONE - Not Robotic (0)
volume group:vault_grp
created:Tue Sep 3 10:08:32 2000
assigned:Tue May 6 00:11:45 2001
last mounted:Tue May 6 11:34:25 2001
first mount:Tue Sep 3 18:20:48 2000
expiration date:---
number of mounts:21
max mounts allowed:---
status:0x1
=====
```

The pool number is listed on the volume pool line, and is the number in between the parentheses. In this case, the media ID would be S04440 and the pool number would be 3.





This chapter covers reporting for the Vault product. Topics in this chapter include:

- ◆ Report Types
- ◆ Reports for Media Going Off Site
- ◆ Reports for Media Coming On Site
- ◆ Detailed Media Reports
- ◆ Recovery Report for Vault
- ◆ Consolidating Reports
- ◆ Running Reports from the Command Line
- ◆ Report Distribution
- ◆ Reprinting Reports

Report Types

This section describes the reports available in Vault. There are three types of reports:

- ◆ Reports for media going off site
- ◆ Reports for media coming on site
- ◆ Detailed media reports

In addition, there is a Recovery Report. We suggest you generate this report on a regular basis, as it is helpful in disaster recovery efforts.

Note In these reports, the tapes used for NetBackup catalog backups do not show a date in the **Assigned** field, but display the notation **NBDBTAPE**. The **Expiration Date** field displays the date calculated as a return date during the assignment. The **#Images** and **KBytes** fields will display zero.



Reports for Media Going Off Site

The *Picking List for Robot* and *Distribution List for Vault* reports show the tapes that have been ejected from the robot and will be transported off site. The *Summary Distribution List for Vault* and the *Detailed Distribution List for Vault* provide more detailed information on the media to be sent off site.

Picking List for Robot

This report is sorted by media ID and should be used by the operations staff as a checklist for media that has been ejected from the robots. You can save the report for tracking purposes, or reprint it as long as the session directory still exists.

Distribution List for Vault

This report is sorted by off-site slot number and should accompany the media that is destined for the off-site vault. The vault vendor should use this report to verify that all the tapes listed were actually received.

Detailed Distribution List for Vault

This report is similar to the *Picking List from Robot* and *Distribution List for Vault* reports except that detailed information is listed for each media. This detail includes the client machines that have backups on this media, when the client was backed up, the NetBackup backup identifier for the backup job and the number of kilobytes stored in this NetBackup fragment.

Keep in mind that backup jobs may span multiple tapes, so it is possible to see duplicate detailed listings of a given backup job on more than one tape. True image recovery (TIR) information listed on this report is indicated by a TIR adjacent to the number of kilobytes. This report is very useful at a disaster recovery site. We recommend you send this report off site.

Summary Distribution List for Vault

This report is similar to the *Detailed Distribution List for Vault* report, except that the entry for each piece of media will list only a unique client, policy, schedule and date. That is, if multiple backup jobs for a given client, policy and schedule (usually seen with RDBMS backups or SAP backups) are written to the same tape on the same date, only one line of information will be printed out on this report. The *Detailed Distribution List* would show each of these backup jobs as a separate entry, which may generate a very long report. The *Summary Distribution List for Vault* report summarizes the information and presents it in a more compact form. This report is also very useful for disaster recovery situations; we recommend you send this report off site.

Reports for Media Coming On Site

The *Picking List for Vault* and *Distribution List for Robot* reports show the tapes that are being requested back from the off-site vault. Vault will not generate these reports until the media have been ejected for the current Vault session.

Picking List for Vault

This report should be sent off site to the vault vendor. Tapes are listed on this report because Vault determined that they are in the appropriate off-site volume group and that they no longer contain valid NetBackup images. When Vault identifies these tapes, it changes the Date Requested field within the Media Manager description field for the media. It then prints out the media ID on this report along with the date requested. Expired media will only appear on the report generated on the date the media expired. If no report is printed for that date, you can access a list of expired off-site media by generating the *Complete Inventory List for Vault* report.

Distribution List for Robot

This report is identical to the *Picking List for Vault*, except that it has a different report title. Retain this report on site to use as a checklist for the media returned from the off-site vault.

Detailed Media Reports

Vault will not generate these reports until the media have been ejected.

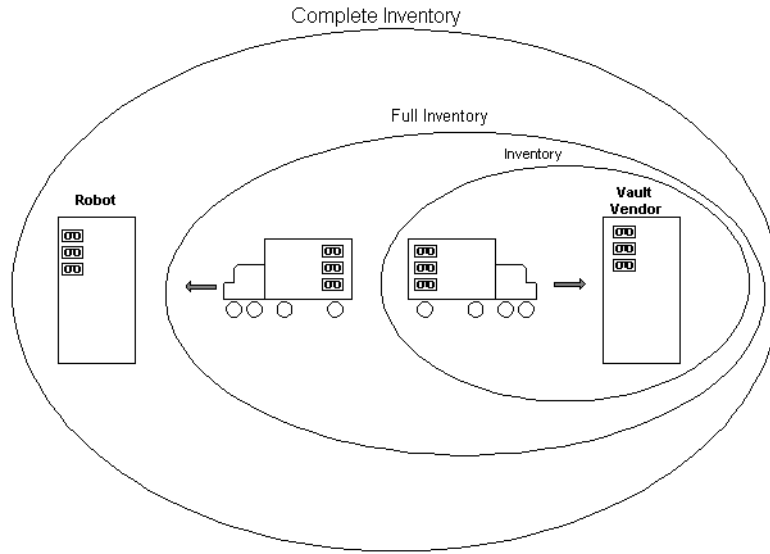
There are three inventory list reports, related in the following way:

- ◆ The Inventory List is the least comprehensive of the three.
- ◆ The Full Inventory List contains everything in the Inventory List Report, and more.
- ◆ The Complete Inventory List contains everything in the Full Inventory List Report, and more.



Graphic Illustration of Scope of Inventory Reports

The different scopes of the reports is demonstrated by the illustration below:



Inventory List for Vault

The *Inventory List for Vault* report shows all media that are currently off site at the vault vendor after the current batch of tapes have been processed and sent off site. This list of tapes is generated by checking the description field for the media, the volume pool and the off-site volume group. If any tapes have been requested to be returned from the vault vendor, they will not be displayed on this report, but on the *Full Inventory List* report. This report is a sanity check for the vault vendor to verify that Vault and the vault vendor agree on the tapes currently off site. We recommend you send this report to your vault vendor to perform this verification.

Full Inventory List for Vault

This report is similar to the *Inventory List for Vault* report except that it includes any tapes that have been requested back from the off-site vault vendor. Normally, this report is not generated on a daily basis. Rather, the *Inventory List for Vault* report is usually sent to the vault vendor to perform verification. The *Full Inventory List for Vault* contains the same information as *Inventory List for Vault* plus information about tapes in transit.

Complete Inventory List for Vault

This report shows all tapes within the off-site volume pool. If the tapes are currently at the vault, the code V is shown in the Location field, while the code R appears for tapes that are currently within the robot. The *Complete Inventory List for Vault* report contains the same information as the *Full Inventory List* for Vault plus information about on-site tapes.

Note Tapes within the off-site volume pool must belong to either the off-site volume group or the Robotic volume group, or they will not appear on this report.

Recovery Report for Vault

This report shows all policies defined on a NetBackup master server and all media that is required for restores between a given set of dates. The title of the report displays the date range covered by the report.

This report includes the NetBackup catalog tapes that are currently off site. For the NetBackup catalog media to be listed in this section, their volume group must match the volume group specified in the off-site volume group. Only NetBackup catalog media that are assigned will appear on this report.

Consolidating Reports

You may want to run your Vault sessions daily and eject tapes at the end of the week. To do this, you would consolidate your ejections and reports as follows:

1. Select **Manual Eject** in the **Eject Mode** area of the **Eject** tab.

This action ensures that tapes will not automatically be ejected for the individual Vault session.

2. Select **Manual Reports** in the **Report Mode** area of the **Reports** tab.

This action ensures that reports will not automatically be generated for the individual Vault session.

Ejecting tapes from multiple sessions can be done manually or as a scheduled event.

- ◆ To do so manually, use `vltopmenu`.
- ◆ To do so as a scheduled event, create a new Vault policy in Policy Management and configure the `vlteject` command as an entry in the file list.



Running Reports from the Command Line

Use `vlteject` to run reports from the command line. Access `vlteject` from the following directory:

UNIX: `/usr/opensv/netbackup/bin`

Windows: `install_path\NetBackup\bin`

`vlteject` is the command which runs `eject` and report requests that have not been run automatically. When you run the `vlteject` command, you can optionally specify a vault, or a vault and a session ID. If you specify neither, the command will run all ejects and reports that have not yet been executed.

The syntax for the command is:

```
vlteject [-auto y/n] [-vault vault_name [-sessionid id]]
        [-eject] [-report ][-eject_delay seconds]
```

where:

vault_name is the name of the vault for which you want to consolidate ejects and/or reports.

id is the number of the session for which you want to run ejects and/or reports.

seconds is the number of seconds to delay the eject.

For a more detailed description of the `vlteject` command, please refer to the “Commands” Appendix.

Option	Description
-eject	If -eject is set, vlteject will run the eject process for all selected vault/session combinations.
-report	If -report is set, vlteject will run all reports specified in the vault/sessions combinations. If the corresponding eject process has been completed, then all enabled reports are generated and distributed. These reports will not be run again if vlteject is run again. If eject has not been completed, then the subset of reports which do not depend on completion of eject will be run. These reports will be run again if vlteject is run again.
-interactive	If -interactive is set, the user can choose where to display the consolidated reports.

Report Distribution

When vlteject is run and ejects and/or reports for more than one session are affected, reports may be collected and distributed in three different ways: by email, to a printer, or consolidated to a file.

Reports from all sessions that have a common destination are concatenated into a single report. If two sessions specify that all reports should be generated, but specify different report destinations, the reports will not be concatenated.

If the profile that a particular session is associated with does not specify a report destination, then reports for that session will not be distributed.

Note You must specify a report destination in the profile for the reports generated to be distributed.

Reprinting Reports

Note You can only reprint reports if the session directory for that vault still exists. Only some of reports will be valid -- for example, the picking list reports are only valid for the current day.

Use the vltopmenu MUI to reprint individual reports. For information on vltopmenu, refer to "Using vltopmenu."





Using the Menu User Interfaces (MUIs)

9

Vault 4.5 has two menu user interfaces. For the most part, we expect you to rely on the GUI to configure and run Vault, however, there are certain functions that can only be achieved through the use of one of Vault's MUIs. There are also many Vault options that you can use either from the GUI or from one of the MUIs.

This chapter covers the following topics:

- ◆ Menu User Interfaces in Vault 4.5
- ◆ Accessing the Menu User Interfaces
- ◆ Help for `vltadm`
- ◆ Using `vltopmenu`
- ◆ Changes in `vmadm` for Vault
- ◆ Changes to `bpdbjobs` for Vault

The chapter may refer to functionality covered in another chapter, but will not go into detail. Refer to the referenced chapter for specific information.

Menu User Interfaces in Vault 4.5

The two MUIs available in Vault 4.5 are:

- ◆ `vltadm`

This menu enables you to configure Vault profiles. It is equivalent to the Vault GUI in terms of functionality.

Note `vltadm` is only available on UNIX systems.

- ◆ `vltopmenu`

This menu provides a way to eject media and print reports for one or more Vault sessions.

At the end of this chapter is a section on changes in the information displayed in Volume Configuration when Vault is in use. You can edit these fields using `vmadm`.



Accessing the Menu User Interfaces

The Vault MUIs reside in the following directory:

UNIX: `/usr/opensv/netbackup/bin/`

Windows: `install_path\NetBackup\bin\`

Using vltadm

Note vltadm is only available on UNIX systems.

vltadm provides a way to configure and run Vault from a text-based menu. The options presented in this menu can also be accessed through the Vault GUI. vltadm is only available to users with root privileges.

This interface can be used from any character-based terminal (or terminal emulation window) for which the administrator has a termcap or terminfo definition. Use vltadm only on the master server and ensure that no other instances of vltadm are active when you are modifying the configuration.

Note Only run one instance of vltadm at one time. Do not run vltadm at the same time as the NetBackup console. Do not run another session for this vault while using this menu.

▼ To run vltadm:

The initial vltadm screen provides a list of the options possible through the MUI. Each of these choices presents a submenu of options. The initial screen is called Vault Administration. For the robot, vault, and profile, **ALL** will be chosen. Use the menu to browse for specific robots, vaults, or profiles. The robot, vault, and profile names at the top of the menu will change. When the one you want to act on displays, simply type the letter of the action you want to perform.

◆ Vault Administration

Robot Name:

Vault Name:

Profile Name:

- r) Browse all configured robots
changes the Robot Name field at the top of the screen
- v) Browse all configured vaults for selected robot
changes the Vault Name field at the top of the screen

- p) Browse all configured profiles for selected vault
changes the Profile Name field at the top of the screen
 - n) Robot Management...
brings up the Robot Management menu
 - t) Vaults for selected robot...
brings up vaults for selected robot
 - f) Profiles for selected vault...
brings up profiles for selected vault
 - c) Copy selected profile...
 - s) Start Session for selected profile....
 - a) Vault Properties...
 - h) Help
brings up the Help Facility menu. See “Help for vltadm” on page 128.
 - q) Quit (without saving)
 - x) Save and Exit
- =====

◆ Vault Properties

You can specify email addresses for notification of session status and default email address for reports. Multiple addresses can be separated by semi-colons(;), or commas (,). If a media server has ever been known by other name or names, those names should be specified as well. Specifying all names by which a media server has been known allows Vault to identify images as belonging to that media server, even if that media server was known by a different name when the image was created.

Notification email:

Reports email:

Media server name:

Alias host name:

- n) Change email address for notification...
- r) Change default email address for reports...
- a) Add media server host name...
- m) Modify media server host name...
- d) Delete media server host name...



- s) Add alias for this host name...
- t) Delete alias for this host name...
- b) Browse media server host names
- h) Help
- p) Previous menu

=====

Choosing **n** from the Vault Administration menu brings up the Robot Management menu:

◆ Robot Management

Use this screen to specify a robot that you want to use for your vault. You can select any robots that are recognized by NetBackup and have storage units on them. Based on the robot number that you select, the other fields will fill automatically.

Robot Number:

Robot Name:

Robot Type:

VolDbHost:

Robotic Control Host:

Output Destination:

- a) New Vault Robot...
- d) Delete Robot...
- s) Save selected robot
- n) Vaults for this robot...
- b) Browse robots forward
- r) Browse robots reverse
- l) List/Display robots...
- o) Output Destination
- h) Help
- p) Previous menu
- x) Save and Exit

=====

Choosing **n** in the Robot Management menu brings up the Vault Management menu.

◆ Vault Management

Robot Name:

Vault Name:

Vault Vendor:

Off-site Volume Group:

Vault Seed:

Robot Volume Group:

Output Destination:

- a) New Vault...
- m) Change Vault...
- d) Delete Vault...
- n) Profiles for this Vault...
- s) Save selected vault
- b) Browse vaults forward
- r) Browse vaults reverse
- l) List/Display vaults...
- o) Output Destination (SCREEN or FILE)
- p) Previous menu
- x) Save and Exit

=====

◆ Profile Management

Robot Name:

Vault Name:

Profile Name:

Duplication step skipped:

Catalog backup step skipped:

Eject step skipped:

Output Destination:

- a) New Profile...
- m) Change Profile...
- d) Delete profile...



- c) Copy selected profile...
- e) Start Session for selected profile....
- s) Save selected profile
- b) Browse profiles forward
- r) Browse profiles reverse
- l) List/Display profiles...
- o) Output Destination (SCREEN or FILE)
- h) Help
- p) Previous menu
- x) Save and Exit

=====

The next five options involve configuration at the profile level.

◆ Choose Backups

Profile Name:

Backups between:

Backup Type:

Source Volume Group:

Clients:

Media Servers:

Backup Policies:

Schedules:

Policy Option:

- t) Change Backups between...
- b) Choose Type of Backup...
- g) Change Source Volume Group...
- c) Change Client list...
- m) Change Media Server list...
- l) Change Backup Policy list...
- s) Change Schedule list...
- a) Change Policy Option

- n) Proceed to next step of profile
- h) Help
- p) Previous menu

=====

◆ Duplication

Profile Name:

Duplication step skipped:

Multiplex while duplicating:

Robot shared by multiple servers:

Expire original disk backups after: hours

Output Destination:

- k) Skip/Enable Duplication
- x) Enable/Disable multiplexing...
- r) Enable/Disable shared robots...
- c) Change expire hours for original disk backups...
- i) Duplication configuration...

brings up the Duplication Items Setup menu

- n) Proceed to next step of profile (Catalog Backups)
brings up the Catalog Backups menu
- l) List/Display duplicate configurations...
- o) Output Destination (SCREEN or FILE)
- h) Help
- p) Previous menu

=====

◆ Duplication Items Setup

Profile Name:

Backup Server:

Alternate Read Host:

Read Drives:

Write Drives:



Primary:

Number of Copies:

- b) Change backup server...
- t) Change alternate read host...
- r) Change read drives...
- w) Change write drives...
- y) Change primary...
- c) Copy configuration...
- a) New duplication item...
- d) Delete this duplication item
- n) Browse the duplication items
- p) Previous menu

=====

◆ Catalog Backup

Profile Name:

Catalog backup skipped:

Media Server:

Catalog Volume Pool:

Retention Period:

Number of (serial) Catalog Backups:

Catalog Backup Policy:

- k) Skip/Enable Catalog Backup
- m) Change Media Server...
- v) Change Catalog Volume Pool...
- r) Change Retention Period...
- c) Specify Number of Catalog Backups to Perform...
- b) Specify Multiple Tape Catalog Backup Policy....
- s) Specify Files to be Backed Up....
- a) Add More Files to Backup List...
- n) Proceed to Next Step of Profile Setup(Eject)...

- h) Help
- p) Previous menu

=====

◆ Eject

Profile Name:

Eject step skipped:

Eject mode:

Suspend media:

Start time:

End time:

- k) Skip/Enable Eject
- m) Change Eject mode
- e) Enable/Disable Suspend media
- t) Change Start time...
- d) Change End time...
- v) Specify Volume Pools to be Ejected...
- n) Proceed to next step of profile setup(Reports)...
- h) Help
- p) Previous menu

Use the Eject screen to configure eject options. Configure the three following fields of information:

◆ Volume pools

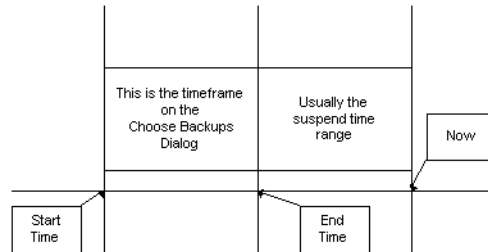
Volume pools from which to eject tapes. Do not put the NetBackup pool into this list. If you are doing Duplication and a Catalog Backup and you want to vault them, then you want to eject both volume pools used for these steps.

◆ Suspend (with time range)

Since images are usually duplicated and ejected in batches, there are some backups that you would like to vault but cannot because they are not complete by the time you want to start vaulting. Suspend those tapes that don't get vaulted; they get picked up in the next batch. (See "Use the Suspend Option to Avoid Vaulting Partial Backups" on page 32.)

The suspend option is the time frame between the *Now point and the *End point, illustrated in the diagram below:





◆ Eject mode

Ejecting tapes automatically or manually. Not all tapes are ejected from the volume pools to be ejected. Only tapes to which duplicated images and catalog backups were written are ejected.

=====

◆ Reports

Use this screen to specify which reports to generate. If you select a manual Report Mode (option m), you must manually generate the reports you have selected. Otherwise Vault will generate most reports automatically. (Some reports cannot be generated until the media is ejected.)

There are three types of reports that you can generate: reports for media coming on site, reports for media going off site, and detailed media reports. For detailed information on these reports, please see “Reporting” on page 109.

Profile Name:

Report mode:

Report Destination File:

Report Destination Printer:

Report Destination Email(s):

- m) Modify Report mode
- f) Specify Report Destination Filename...
- t) Specify Report Destination Printer...
- e) Specify Report Destination Email(s)...

- r) Specify reports to be generated...
- c) Change report titles
- d) Change report header...
- h) Help
- p) Previous menu



Help for vltadm

The help facility provides the user with online text explaining the various features of vltadm. To enter the help facility choose the Help (h) option on the main menu or any submenus.

- m) Help on vltadm initial menu
displays information about the initial menu of vltadm, with brief descriptions of the individual submenus
- t) Tutorial
introduces vltadm, with a brief description of Vault concepts
- h) Help (on Help)
displays this information
- p) Previous menu
returns to the previous menu

Choosing the help option in any other menu will display help text for the current menu.

How to Use the Help Screens

At the bottom of the help screens are the controls to maneuver around the display. The bottom line of the display looks like this:

(B)ack (F)orward (U)p (D)own (Q)uit

Pressing the key in parentheses causes that action to be taken. The keys are not case sensitive. These actions are described as follows:

- **Back** scrolls the display up one screen of text (unless the display is at the top of the text).
- **Forward** scrolls the display down one screen.
- **Up** scrolls the display up one line.
- **Down** scrolls the display down one line.
- **Quit** returns display to the previous menu.

Using vltopmenu

vltopmenu is the Vault operations menu. It allows you to eject tapes and print reports for one or more Vault sessions. For information about each report, please refer to “Reporting” on page 109.

Upon start up, the menu will show no profile. You need to select the profile or the *robot_number/vault_name/profile* parameter before you begin to use the menu. If you do not change the vault or the session specified, these vltopmenu options will apply to the most recent session.

You can view the results of each operation you execute with the menu in a log file. The name and location of the log file is located at the end of the output for each command.

For example, if you choose **Eject media for this session [3]** on a UNIX system, the output is:

```
vlteject Started
```

```
vlteject Completed
```

The results of this operation have been logged in the following file:

```
/usr/opensv/netbackup/vault/sessions/vlteject_status/  
details.log.timestamp
```

Note Do not run another session for this vault while using this menu.

- ❖ The initial vltopmenu screen displays the current profile, session, and path for the Print command followed by the list of options. Type the number of the option you want to perform after the word **Selection** at the bottom of the menu.

The options are detailed below.

- ◆ Change Profile [1]

This option allows the user to perform the actions allowed in this menu on any profile, one-at-a-time, without closing the application.

- ◆ Change Session [2]

You may need to run vltopmenu for a session other than the default session if you are trying to eject or report for a particular vault session.

- ◆ Eject Media for This Session [3]

This option is similar to the ACSLS `eject media` command except that it generates CLI calls which automatically eject tapes from the robot.



Note The tapes to be ejected must belong to the appropriate robotic volume group and duplication pool, and must be assigned a unique off-site slot number.

◆ Run Reports for This Session [4]

This step prints all the reports that are generated as part of this session (identified by the sessionID and profile). The reports will be sent using the print command.

◆ Run Individual Reports [5]

This step lists all the reports so that you can choose the reports to run. It also includes a Change Print Command option. The reports will be sent using the print command, unless you change the command.

◆ Inject Tapes into Robot [6]

When tape inject is occurring, the volume group changes to the robotic volume group.

◆ Consolidate All Reports [7]

This option generates reports for any vault that has not had reports generated for a given session.

◆ Consolidate All Ejects [8]

This option ejects media for any vault that has not had media ejected for a given session.

◆ Consolidate All Reports and Ejects [9]

Consolidate ejects media from vault sessions and runs the corresponding reports as selected in a profile. Consolidate can:

- Eject media and run reports for a single session as designated by vault and session parameters
- Eject media and run reports for the sessions within a vault for which these actions have not occurred.
- Eject media and run reports for the sessions within all vaults for which these actions have not occurred.

◆ Exit [q]

Use this option to quit the interface.

Changes in vmadm for Vault

When Vault is in use, several fields are added to the Volume Configuration display. You can modify these fields through the Special Ops menu in vmadm.

Additions to Volume Configuration

When you display the full Volume Configuration the following information for Vault is displayed. For example,

volume group:TL8-0

vault name: V1

vault sent date: Wed Dec 02 09:34:01 2000

vault return date: Tue Feb 17 09:34:01 2001

vault slot: 546

vault session ID: 37

created: Mon Nov 29 08:29:03 2000

Changes to the Special Actions Menu

There are several Vault options that you can modify through the Special Ops menu in vmadm. These are described below.

Change Vault Name for Volumes

You can set, clear, or change the name of the vault that the volume is contained in. This field is used by Vault to determine where the volume is located when it is at an off-site location.

Caution If you are changing vault names for volumes and you used the legacy product, you will have problems!

▼ To modify the name of the vault:

1. On the main menu, choose **s** for **Special Actions**.
2. Choose **a** for **Change Vault Parameters for Volumes**.
3. Choose **n** for **Change Vault Name for Volumes**.



4. Enter the new vault name (25 character maximum). Enter a hyphen (-) to clear the field.
5. You will be prompted for the media IDs for which you want this vault name applied. The prompt will repeat until you press the enter key without entering a media ID. Click the ESC key to cancel the action.

Change Date Volumes are Sent to Vault

You can set, clear, or change the date a volume is sent to the off-site vault. This field is used by Vault to record when a volume was sent to the off-site vault location. You can modify this date for a single volume or for multiple volumes.

▼ To modify the Vault Sent Date of the vault:

1. On the main menu, choose **s** for **Special Actions**.
2. Choose **a** for **Change Vault Parameters for Volumes**.
3. Choose **d** for **Change Date Volumes are Sent to Vault**.
4. Enter the new date the volume was sent off-site. Enter a zero (0) to clear the field.
5. You will be prompted for the dates for which you want this vault name applied. The prompt will repeat until you press the enter key without entering a date. Click the ESC key to cancel the action.

Change Date Volumes Return from Vault

You can set, clear, or change the date a volume returns from the off-site vault. This field is used by Vault to record when a volume is requested to return from the off-site vault location. You can modify this date for a single volume or for multiple volumes.

▼ To modify the Vault Return Date of the vault:

1. On the main menu, choose **s** for **Special Actions**.
2. Choose **a** for **Change Vault Parameters for Volumes**.
3. Choose **r** for **Change Date Volumes Return from Vault**.
4. Enter the new date the volume is requested to return from the off-site vault. Enter a zero (0) to clear the field.

5. You will be prompted for the return dates for which you want this vault name applied. The prompt will repeat until you press the enter key without entering a return date. Click the ESC key to cancel the action.

Change Vault Slot for Volumes

You can set, clear, or change the slot that the volume is contained in at the off-site location. This field is used by Vault to determine in what slot the volume is located in the off-site vault. You can modify the slot for a single volume or for multiple volumes.

▼ To modify the Slot ID for a volume:

1. On the main menu, choose **s** for **Special Actions**.
2. Choose **a** for **Change Vault Parameters for Volumes**.
3. Choose **s** for **Change Vault Slot for Volumes**.
4. Enter the new vault slot ID for the volume. Enter a zero (0) to clear the field.
5. You will be prompted for the slot IDs for which you want this vault name applied. The prompt will repeat until you press the enter key without entering a slot ID. Click the ESC key to cancel the action.

Change Vault Session ID for Volumes

You can set, clear, or change the session ID in which a volume was processed. This field is used by Vault to determine in what vault session a volume was processed. You can modify the session ID for a single volume or for multiple volumes.

▼ To modify the Session ID for a volume:

1. On the main menu, choose **s** for **Special Actions**.
2. Choose **a** for **Change Vault Parameters for Volumes**.
3. Choose **i** for **Change Vault Session ID for Volumes**.
4. Enter the new session ID for the volume. Enter a zero (0) to clear the field.
5. You will be prompted for the session IDs for which you want this vault name applied. The prompt will repeat until you press the enter key without entering a session ID. Click the ESC key to cancel the action.



Changes to bpdbjobs for Vault

The Activity Monitor MUI, `bpdbjobs`, displays Vault jobs by default, along with NetBackup jobs. If the `-vault` option is used:

```
bpdbjobs -vault
```

then the display includes the Vault job fields described below.

Field	Description
Robot	The name of the robot the vault is associated with.
Vault	The name of the vault under which this session is running.
Profile	The name of the profile that holds the configuration information for this vault session.
Session ID	This is the session ID, a unique numeric value, for this vault job. Session ID assignment starts at 1 the first time a vault job is run after vault has been installed. The value is incremented by one every time a new vault job runs.
Tapes to Eject	The number of tapes to be ejected for a vault session. Note If the profile is configured for manual eject, the tapes may not have been ejected yet.
Operation	For vault jobs, the field contains one of the following values. These values progress from the first value to the last as the vault job progresses: <ul style="list-style-type: none"> ♦ Choosing Images ♦ Duplicating Images ♦ Choosing Media ♦ Catalog Backup ♦ Eject and Report ♦ Done

If a vault job completes successfully (with exit status = 0), then the State field and the Operation field both contain the value Done. If a vault job fails, then the Operation field contains the operation that was happening at the time the job failed.

This chapter discusses potential problems that can occur when using Vault and how to resolve or work around them. The discussion is divided into the following parts:

- ◆ Errors Returned by the Vault Session
- ◆ Other Troubleshooting Issues
- ◆ Logs To Accompany Problem Reports

Errors Returned by the Vault Session

Every Vault session writes a detailed error status to `stderr`.

- ◆ If the error generated by the Vault session is less than or equal to 255, it will return the actual error code. Error codes less than or equal to 255 (except 252) map to standard NetBackup error codes and are documented in the *NetBackup Troubleshooting Guide*.
- ◆ If the Vault session fails with an error code greater than 255, it will return error code 252 and the actual error code will be written to `stderr`. This is because codes greater than 255 are called NetBackup Extended Error Codes and are not supported by all operating systems.

The format of the error text written to `stderr` is:

EXIT status = *error code*

For detailed information on status codes, see the *NetBackup Troubleshooting Guide*.

Other Troubleshooting Issues

Media Missing in Robot

Duplication may fail if a requested piece of media is not found in the robot. The options **Use Inventory to Update Volume Configuration** and **Robot Inventory** within the Media Manager GUI can be used to compare the tapes actually stored in the robot with the



Media Manager database. You must fix the problem by determining where the tape is actually stored. Use the Media Manager GUI to move the tape once it has been found. If the tape is in the vault, the volume group should be set to the vault group.

If the tape is not found, you should delete it from the NetBackup system. If the tape is missing yet is assigned and has valid duplicate images, you will need to use the command `bpexpdate`, which is documented in the *NetBackup System Administration Guide*, to expire the images before you delete the tape from Media Manager.

Bad or Missing Duplicate Tape

If a duplicate tape is lost or damaged, you can reduplicate images found on the tape. The steps required are:

1. Determine which images were on the tape. See the appendix for this procedure.
2. Expire the lost or damaged duplicate tapes using the command `bpexpdate`.
3. Determine when the images were originally created. Use the command `bpimagelist -U -backupid backupid` to see this date.
4. Create a profile containing policy entries only for the policy names found for these images, and set the `duplicate_days` to be greater than the days shown above. For example, set it to 32 if the primary copy was made one month ago.

Note You can run duplicates again only if you have a primary copy.

Make sure that the normal Vault session has finished running for the day before running this second instance. Also, you can double-check that the lost duplicate images were stored in the `preview.list` file (located in the working directory for this duplication session) before proceeding with **Run Duplicates**.

5. Use the Report tab in the Vault interface to print a set of reports for this run. Be sure to provide both sets of reports to the vault vendor.

If You Need to Stop Vault

▼ To stop a Vault session:

1. From the NetBackup Activity Monitor, highlight the Vault job you want to stop.
2. From the **Actions** menu, select **Cancel Job**.

Tape Drive or Robot Offline

If you have a problem with ACSLS drives going offline. These problems will show up on the Media Manager GUI. You should normally try UP'ing the drive, resetting the drive or both. If drives persistently go offline, duplication may hang. There may be a problem with keeping track of how many drives are available when duplicates are running.

Also, if the tape drives are listed as AVR control in the Device Management GUI, there may be a problem with the robotics control. All drives should normally be listed as robotically controlled (e.g. TLD, ACS, etc.), but they will be converted to AVR control if a problem occurs with the robot. There should be an error message in the system logs (for example, `/var/adm/messages`) that will help you diagnose the cause of this problem. You can also use the robotic test utilities (e.g. `robtest`) to help further debug the problem.

No Duplicate Progress Message

If you see a message similar to the following example in the vault debug log `detail.log`, then the Vault process has not received any new information from the `bpduplicate` process within the time frame specified (in this case, 30 minutes):

```
bpduplicate_progress_logname: no activity in 30 minutes
```

bpduplicate_progress_logname is the progress log that `bpduplicate` creates as it runs the duplication for Vault. This file resides in `vault_session_directory/duplicate.log.n`

where *n* is 1 for the first instance of `bpduplicate` that the Vault session executes, 2 for the second instance in the same session, and so on. Its value is incremented by one for each additional instance of `bpduplicate` run during the vault session.

This message does not necessarily indicate an error has occurred. If the image that is currently being duplicated is very large, for example, several gigabytes, then this message is displayed only for informational purposes. To determine if this represents a problem, you may wish to find out the size of the current image. Look at the last few lines of the file `details.log`. This will tell you the backup image ID. With this information, execute the following command, modified to specify the actual image identifier for your Vault session's current image. (The `bpimagelist` command is documented in the *NetBackup System Administrator's Guide*)

```
UNIX: # bpimagelist -L -backupid server2_0897273363
```

```
Windows: C:\Veritas\NetBackup\bin\admincmd bpimagelist.exe -L  
-backupid server2_0897273363
```

The output of this command will show you various statistics about this backup image, including the number of kilobytes written during this backup. If the number is relatively small, there may be a problem with the duplication process. Sometimes this delay is caused by a media mount (which normally does not occur in robotic devices during



duplication) or hardware problems, or the media might be in use. Examine the Device Management interface to determine if there are any hardware problems and also check the system logs. If the backup image is very large, then this message should be regarded as informational.

Ejecting Tapes While in Use

If Vault is configured to eject original media, it is possible that a piece of media could be in use during the eject process (for example, for a restore, or a media verify procedure.) In this case, an error message will be generated by Media Manager. A similar error may be generated by non-Media Manager controlled robots if a piece of media is currently in use.

If you receive one of these errors, we recommend that you use the operations menu `vltopmenu` to re-eject the media after the media is no longer in use. You may receive additional error messages because the rest of the media for the scheduled job has already been ejected.

Ejecting More Media Than Export Capacity

If automatic eject is specified in the profile and a Vault session produces more media for eject than will fit in the media access port (MAP), then the Vault job will not perform an automatic eject. After the job finishes, the operator will need to manually eject the media and generate the reports. Manual eject and manual reporting must be done using `vltopmenu` or `vlteject`.

Vault Session Locking

For a given vault name, only one vault session can run at a time. If a second vault session is started while a first is already active for the vault name, the second vault session will immediately terminate. These sessions may have been started by one or more profiles within a single vault. If this locking situation is encountered, the second vault session will immediately terminate with the message:

A session is already running for this vault

The vault lock file is located in the directory for the specified vault:

UNIX: `/usr/opensv/netbackup/vault/sessions/vault_name/vault.lock`

Windows:

`install_path\NetBackup\vault\sessions\vault_name\vault.lock`

Logs To Accompany Problem Reports

In order to debug the problems, the problem report must be accompanied by a complete set of logs. If the debug logs have not been turned on, you may need to rerun the test by selecting a debug level of 5 for vault and after creating the following debug directories on the master server:

UNIX:

```
/usr/opensv/netbackup/vault/logs/vault
/usr/opensv/netbackup/vault/logs/bpbmvlr
/usr/opensv/netbackup/vault/logs/bpcdr
/usr/opensv/netbackup/vault/logs/admin
/usr/opensv/netbackup/vault/logs/bpsched
```

Windows:

```
install_path\NetBackup\logs\vault
install_path\NetBackup\logs\bpbmvlr
install_path\NetBackup\logs\bpcdr
install_path\NetBackup\logs\admin
install_path\NetBackup\logs\bpsched
```

For problems related to reading or writing of images on the Media server, the following directories may need to be created on the Media Server:

UNIX: `/usr/opensv/netbackup/logs/bptm`

Windows: `install_path\NetBackup\logs\bptm`

For problems related to vault jobs, please rerun the job by adding `-verbose` option to the `vltrun` command added to the file list in the NetBackup policy of type vault.

Alternatively, you can rerun the vault job after adding the following directive to the `bp.conf` file:

```
VAULT_VERBOSE = 5
```

Note that the above directive will impact all the vault jobs. Whereas the `-verbose` option applies only to the vault job that is started via a command that contains the option.

A problem report must be accompanied with the following logs:

- debug log from:

```
UNIX: /usr/opensv/netbackup/logs/vault/log.mmdyyr
```



- Windows: *install_path\NetBackup\logs\vault\mmddyy.log*
- debug log from:
UNIX: */usr/opensv/netbackup/logs/bpcd/log.mmddyy*
Windows: *install_path\NetBackup\logs\bpcd\mmddyy.log*
 - debug log from:
UNIX: */usr/opensv/netbackup/logs/bpbrmvlr/log.mmddyy*
Windows: *install_path\NetBackup\logs\bpbmvlr\mmddyy.log*
 - debug log from:
UNIX: */usr/opensv/netbackup/logs/admin/log.mmddyy*
Windows: *install_path\NetBackup\logs\admin\mmddyy.log*
 - The session.last file in the following directory:
UNIX: */usr/opensv/netbackup/vault/sessions/vault_name*
Windows: *install_path\NetBackup\vault\sessions\vault_name*
 - Contents of the session directory:
UNIX: */usr/opensv/netbackup/vault/sessions/vault_name/sidxxx*
Windows:
install_path\NetBackup\vault\sessions\vault_name\sidxxx

For consolidated eject/reports via vlteject or vltopmenu, see:

UNIX: */usr/opensv/netbackup/vault/sessions/vlteject.mstr*
Windows: *install_path\NetBackup\vault\sessions\vlteject.mstr*

In some cases the following logs on the media servers may be of interest:

UNIX: */usr/opensv/netbackup/logs/bptm/log.mmddyy*
Windows: *install_path\NetBackup\logs\bptm\mmddyy.log*

These commands can be run from the command line only. They are not available through the interface. You must be in the *install_path/bin* directory.

On the following pages, you will find a description of the following commands:

- ◆ vltadm
- ◆ vlteject
- ◆ vltinject
- ◆ vltoffsitemedia
- ◆ vltopmenu
- ◆ vltrun

To run a vault session from the command line, set your environment variable `PATH` to contain the path in which the NetBackup binaries are installed. You must specify the robot number, vault, and profile; or you can specify just the profile, if it has a unique name. For example,:

```
vltrun 1/your_vault/your_profile
```

where *1* is the robot number, *your_vault* is vault, and *your_profile* is the profile.

User Requirements

`vltinject`, `vlteject`, and `vltadm` must be run by root users.

`vltoffsitemedia`, `vltopmenu`, and `vltrun` can be run by non-root users.



vltadm

NAME

vltadm - Start the NetBackup Vault menu interface for administrators.

SYNOPSIS

```
/usr/opensv/netbackup/bin/vltadm
```

DESCRIPTION

The vltadm utility is a menu interface that an administrator can use to configure NetBackup Vault. Running vltadm requires root privileges. This interface can be used from any character-based terminal (or terminal emulation window) for which the administrator has a termcap or terminfo definition.

See the NetBackup Vault *System Administrator's Guide* and the vltadm online help for detailed operating instructions.

FILES

UNIX:

```
/usr/opensv/netbackup/help/vltadm/*  
/usr/opensv/netbackup/db/vault/vault.xml  
/tmp/bp_robots  
/tmp/bp_robots  
/tmp/bp_vaults  
/tmp/bp_profiles  
/tmp/bp_duplicates  
/tmp/_tmp
```

Windows:

```
install_path\NetBackup\help\vltadm  
install_path \NetBackup\db\vault\vault.xml  
\\temp\\bp_robots  
\\temp\\bp_robots  
\\temp\\bp_vaults
```

```
\temp\bp_profiles  
\temp\bp_duplicates  
\temp\_tmp
```

EXIT STATUS

Vault may exit with a status code greater than 255. Such status codes are called extended exit status codes. For such a case, the exit status returned to the system is 252, and the actual exit status is written to `stderr` in the format, `EXIT status = exit status`

The extended exit status values are documented in the *NetBackup Troubleshooting Guide* and in the *NetBackup Troubleshooting Wizard*.

SEE ALSO

```
vltrun(1m)
```



vlteject

NAME

vlteject - Eject media and/or generate reports for previously run sessions

SYNOPSIS

The syntax for the command is:

```
vlteject -eject [-vault vault_name [-sessionid id]][-auto y|n]  
          [-eject_delay seconds]  
  
vlteject -report [-vault vault_name [-sessionid id]]  
  
vlteject -eject -report [-vault vault_name [-sessionid id]]  
          [-auto y|n] [-eject_delay seconds]
```

where auto implies non-interactive, n is default (interactive).

DESCRIPTION

vlteject ejects media and generates the corresponding reports (as configured in the profiles) for vault sessions for which media have not yet been ejected. vlteject can process the pending ejects and/or reports for all sessions pertaining to a single vault or for all sessions pertaining to all vaults.

Depending on how it is called it can run interactively or not. Running interactively is most useful when you will be ejecting more media than will fit in the media access port.

vlteject operates only on sessions for which the session directory still exists. After that directory is cleaned up (removed by NetBackup) vlteject can no longer eject or report for that session.

If you create a directory named:

UNIX: /usr/opensv/netbackup/logs/vault

Windows: *install_path*\netbackup\logs\vault

with public-write access, vlteject will create a daily debug log called log.DDMMYY (where DDMMYY is the current date) file in this directory that can be used for troubleshooting. Public-write access is needed because not all executables that write to this file run as root. The host property **Keep vault logs for n days** determines how long the vault session directories are retained.

Do not modify your vault configuration while vlteject is running.

vlteject can be run in any of the following ways:

- ◆ Directly from the command line

- ◆ By NetBackup policy scheduling. In this case, the policy must be of type Vault, and the policy's file list must consist of a `vlteject` command.
- ◆ Using `vltopmenu` to run an eject operation or a consolidated eject or consolidated report operation

To do a consolidated eject for all robots, do not specify the `-vault` option.

OPTIONS

- `-auto y|n`
`y` (the default) causes `vlteject` to use the options provided on the command line, rather than running interactively. `n` causes `vlteject` to run interactively.
- `-vault vault_name`
The vault name. If the `-vault` option is not specified, `vlteject` will operate on all sessions for all vaults.
- `-sessionid id`
The numeric session id. If the `-vault` option is specified, but the `-sessionid` option is not specified, `vlteject` will operate on all sessions for the specified vault.
- `-eject`
Indicates that the caller wishes eject to be attempted for the indicated sessions. This is optional in case eject has been done and only printing is desired. The eject process must be completed for a session before it is possible to run reports via `vlteject`.
- `-report`
Indicates that the caller wishes to run reports for the indicated sessions. Reports will only be run if eject has been done for each session.
- `-eject_delay seconds`
The number of seconds to delay before ejecting. This is desirable if an operation such as backing up or duplication has just occurred on the affected media. The default is 0. The maximum is 3600 (one hour).
- `-help`
Displays a synopsis of command usage when it is the only option on the command line.

EXAMPLES

To eject media and generate reports for all robots having sessions for which media have not yet been ejected, enter the following:

```
vlteject -eject -report
```



To eject all media that has not yet been ejected for all sessions for the CustomerDB vault, and to generate corresponding reports, enter the following:

```
vlteject -vault CustomerDB -eject -report
```

FILES

UNIX:

```
/usr/opencv/netbackup/db/vault/vault.xml  
/usr/opencv/netbackup/logs/bpbrmvlt/log.mmddyy  
/usr/opencv/netbackup/logs/vault/log.mmddyy  
/usr/opencv/netbackup/vault/sessions/vault_name/sidxxx/  
detail.log  
/usr/opencv/netbackup/vault/sessions/vault_name/sidxxx/  
summary.log  
/usr/opencv/netbackup/vault/sessions/vault_name/sidxxx/  
vlteject_status  
/usr/opencv/netbackup/vault/sessions/vlteject.mstr  
/usr/opencv/netbackup/bp.conf
```

Windows:

```
install_path\NetBackup\db\vault\vault.xml  
install_path\NetBackup\logs\bpbrmvlt\mmddyy.log  
install_path\NetBackup\logs\vault\mmddyy.log  
install_path\NetBackup\vault\sessions\vault_name\sidxxx\  
detail.log  
install_path\NetBackup\vault\sessions\vault_name\sidxxx\  
summary.log  
install_path\NetBackup\vault\sessions\vault_name\sidxxx\  
vlteject.status  
install_path\NetBackup\vault\sessions\vlteject.mstr  
install_path\NetBackup\bp.conf
```

EXIT STATUS

Vault may exit with a status code greater than 255. Such status codes are called extended exit status codes. For such a case, the exit status returned to the system is 252, and the actual exit status is written to `stderr` in the format, `EXIT status = exit status`

The extended exit status values are documented in the *NetBackup Troubleshooting Guide* and in the *NetBackup Troubleshooting Wizard*.

SEE ALSO

`vltopmenu(1)`



vltinject

NAME

`vltinject` – inject volumes into a robot for a specified vault configuration

SYNOPSIS

```
vltinject profile | robot/vault/profile [-f new media_id file] [-if inventory filter value]
```

DESCRIPTION

`vltinject` injects volumes into a robot and updates the Media Manager volume database. It accomplishes this by running the `vmupdate` command, giving it the robot number, robot type, and robotic volume group from the vault configuration matching the specified profile.

If you create a directory named:

UNIX: `/usr/opencv/netbackup/logs/vault`

Windows: `install_path\netbackup\logs\vault`

with public-write access, `vltinject` will create a daily debug log called `log.DDMMYY` (where `DDMMYY` is the current date) file in this directory that can be used for troubleshooting. Public-write access is needed because not all executables that write to this file run as root.

OPTIONS

profile | *robot/vault/profile*

The name of a profile or a robot number, vault, and profile nested within the vault configuration file. If *profile* is used without *robot* and *vault*, the profile must be unique. `vltinject` executes `vmupdate` with the robot number, robot type, and robotic volume group from this profile's configuration.

EXAMPLE

To inject volumes that were vaulted by the Payroll profile and that have been returned from the offsite vault, the user would enter the following:

```
% vltinject Payroll
```

To inject volumes that were vaulted by the Weekly profile in the Finance vault and that have been returned from the offsite vault, the user would enter the following:

```
% vltinject 8/Finance/Weekly
```

To inject volumes that were vaulted by the Payroll profile and that have been returned from the offsite vault, enter the following:

```
% vltinject Payroll
```

To inject volumes that were vaulted by the Weekly profile in the Finance vault, and that have been returned from the offsite vault, enter the following:

```
% vltinject 8/Finance/Weekly
```

FILES

UNIX: `/usr/opensv/netbackup/logs/vault/log.mmddyy`

Windows: `install_path\NetBackup\logs\vault\mmddyy.log`

EXIT STATUS

0 The Volume Database was successfully updated.

>0 There was a problem updating the Volume Database.

Vault may exit with a status code greater than 255. Such status codes are called extended exit status codes. For such a case, the exit status returned to the system is 252, and the actual exit status is written to `stderr` in the format, EXIT status = *exit status*

The extended exit status values are documented in the *NetBackup Troubleshooting Guide* and in the *NetBackup Troubleshooting Wizard*.



vltoffsitemedia

NAME

`vltoffsitemedia` - list the offsite parameter values for a group of media, or change the offsite parameter value for a single media.

SYNOPSIS

```
vltoffsitemedia -list [-vault vault_name] [-voldbhost  
    host_name]  
  
vltoffsitemedia -change -m media_id [-voldbhost host_name]  
    [-vltname vault_name] [-vltsent mm/dd/yyyy  
    [hh:mm:ss]][-vltreturn mm/dd/yyyy [hh:mm:ss]]  
    [-vltslot slot_no] [-vltsession session_id]
```

DESCRIPTION

Allows the user to change the vault-specific parameters of a given media. This allows the user to change one or more parameters using a single command. It allows the user to view the various vault parameters of all media for a particular volume database host or vault.

If you create a directory named:

UNIX: `/usr/opensv/netbackup/logs/vault`

Windows: `install_path\netbackup\logs\vault`

with public-write access, `vltoffsitemedia` will create a daily debug log called `log.DDMMYY` (where `DDMMYY` is the current date) file in this directory that can be used for troubleshooting. Public-write access is needed because not all executables that write to this file run as root.

OPTIONS

`-vault vault_name`
Name of the vault for which all media ids and their vault-specific parameters are to be listed.

`-voldbhost host_name`
Name of the volume database host.

`-m media_id`
Media ID of the media whose vault parameters are to be changed.

`-vltname vault_name`
Specifies the name of the logical vault configured for the robot that ejected the volume.

- `-vltsent mm/dd/yyyy hh:mm:ss`
Specifies the date and time the media was sent to the offsite vault.
- `-vlreturn mm/dd/yyyy hh:mm:ss`
Specifies the date and time the media was requested for return from the vault vendor. For Catalog Backup volumes, this is the date that the media will be requested for return from the vault vendor.
- `-vltslot slot_no`
Specifies the vault vendor's slot number for the slot that this volume occupies.
- `-vltsession session_id`
Specifies the identifier of the Vault session that ejected this media.

EXAMPLES

The following command will change the vault name and the vault sent dates of the media with the ID BYQ123:

```
vltoffsitemedia -change -m BYQ123 -vltname THISTLE
-vltsent 08/01/2001 12:22:00
```

The following command will change the vault slot number to 100 for a media with ID 000012:

```
vltoffsitemedia -change -m 000012 -vltslot 100
```

The following command can be used to clear out the vault-specific fields for a media:

```
vltoffsitemedia -change -m 000012 -vltname ""
-vltsession 0 -vltslot 0 -vltsent 0 -vlreturn 0
```

or:

```
vltoffsitemedia -change -m 000012 -vltname -
-vltsession 0 -vltslot 0 -vltsent 00/00/00 -vlreturn
00/00/00
```

The `vltoffsitemedia` command uses the media manager commands to query/update the volume database. If the `vltoffsitemedia` command fails, look at the debug log in the `install_path/netbackup/logs/vault` directory for detailed information about the actual media manager command that failed. Status codes returned by media manager commands are documented in Chapter 5 of the *NetBackup Troubleshooting Guide*, Media Manager Status Codes and Messages.



EXIT STATUS

Vault may exit with a status code greater than 255. Such status codes are called "extended exit status codes". For such a case, the exit status returned to the system is 252, and the actual exit status is written to `stderr` in the format, `EXIT status = exit status`

The extended exit status values are documented in the *NetBackup Troubleshooting Guide* and in the *NetBackup Troubleshooting Wizard*.

vltopmenu

NAME

vltopmenu - Start the NetBackup Vault menu interface for operators

SYNOPSIS

```
/usr/opensv/netbackup/bin/vltopmenu
```

DESCRIPTION

Allows the user to invoke a menu screen containing the various options that an Operator of the NetBackup Vault feature can use. It allows the user to eject or inject media, print various reports individually or collectively, as well as consolidate all reports and ejects for all sessions which have not yet ejected media. This interface can be used from any character-based terminal (or terminal emulation window) for which the user has a termcap or terminfo definition.

See the *NetBackup Operator's Guide* for detailed operating instructions.

FILES

UNIX:

```
/usr/opensv/netbackup/vault/sessions/vlteject.mstr  
/usr/opensv/netbackup/vault/sessions/vlteject_status.log.  
timestamp  
/usr/opensv/netbackup/vault/sessions/*/sid*/detail.log
```

Windows:

```
install_path\NetBackup\vault\sessions\vlteject.mstr  
install_path\NetBackup\vault\sessions\vlteject_status.log.  
timestamp  
install_path\NetBackup\vault\sessions\*\sid*\detail.log
```

EXIT STATUS

Vault may exit with a status code greater than 255. Such status codes are called extended exit status codes. For such a case, the exit status returned to the system is 252, and the actual exit status is written to `stderr` in the format, EXIT status = *exit status*

The extended exit status values are documented in the *NetBackup Troubleshooting Guide* and in the *NetBackup Troubleshooting Wizard*.



vltrun

NAME

vltrun - Run a NetBackup Vault session

SYNOPSIS

```
vltrun profile | robot\vault\profile | robot/vault/profile
      [-preview] [-verbose|-v]
```

DESCRIPTION

vltrun drives a NetBackup Vault session by issuing a sequence of calls to the vault engine. Optionally, the session can include callouts to user-provided notify scripts.

OPTIONS

profile | *robot/vault/profile*

The name of a profile or a nested robot number, vault, and profile in the vault parameter file. If *profile* is used without *robot* and *vault*, the profile must be unique within the vault parameter file. This option is required.

[-preview]

Generate the Preview list of images to be vaulted in a vault session. The results go to the file `preview.list` in the session directory..

[-verbose|-v]

Report verbosely on the session in the vault debug log.

-help

Displays a synopsis of command usage when it is the only option on the command line.

USAGE

The vltrun session follows this sequence:

- ◆ Run the `vlt_start_notify` script
- ◆ Inventory media
- ◆ Initialize Media Manager database for vault media returned to the robot
- ◆ Generate the list of preview images to be vaulted
- ◆ Duplicate images
- ◆ Inventory Media Manager database (first time)
- ◆ Assign media for the NetBackup catalog backup

- ◆ Inventory Media Manager database (second time)
- ◆ Inventory images
- ◆ Suspend media
- ◆ Run the `vlt_end_notify` script
- ◆ Re-inventory images
- ◆ Assign slot IDs
- ◆ Backup the NetBackup catalog
- ◆ Inventory the Media Manager database (third and final time)
- ◆ Generate the eject list
- ◆ Eject and report
- ◆ Run the `vlt_end_notify` script

`vltrun` can be run in any of the following ways:

- ◆ directly from the command line;
- ◆ by NetBackup policy scheduling. In this case, the policy must consist of type Vault, and the policy's file list must consist of a `vltrun` command;
- ◆ by running the command `Start Session` for a profile in the Vault GUI or `vltadm`.

`vltrun` uses the option `profile | robot/vault/profile` to run a vault session. You can use the `profile` form of the option if there is no other profile with the same name in your vault configuration. In this case, the profile name is sufficient to uniquely identify the configuration information.

If there is more than one profile with the same name, then use the `robot/vault/profile` form to uniquely identify the configuration.

Do not modify your vault configuration while a vault session is running.

When the session starts, it creates a directory to hold the files created by `vltrun` and the vault engine during the session.

The vault session directory is:

UNIX: `/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx`

Windows: `install_path\NetBackup\vault\sessions\vault_name\sidxxx`

`xxx` is an integer uniquely assigned to this session. For each vault name, session identifiers are sequentially assigned, starting with 1.

If you have configured an email address in your vault properties, then email will be sent to this address at the end of the session, reporting the results. By default, email is sent to root.



`vltrun` produces an overview of the session, called `vltrun.log`, in the session directory.

You can control vault processing at several points in the session by installing notify scripts in the directory for NetBackup binaries, `/usr/opensv/netbackup/bin`. Refer to the *NetBackup Vault System Administrator's Guide* for more information on notify scripts.

You can monitor the progress of your `vltrun` session in the NetBackup Activity Monitor. The Operation field on the main Activity Monitor window shows the progress of your vault session:

- ◆ Choosing Images
- ◆ Duplicating Images
- ◆ Choosing Media
- ◆ Catalog Backup
- ◆ Eject and Report
- ◆ Done

If you create a directory named:

(UNIX) `/usr/opensv/netbackup/logs/vault`

(Windows) `install_path\NetBackup\logs\vault`

with public-write access, `vltrun` will create a daily debug log called `log.DDMMYY` (where `DDMMYY` is the current date) file in this directory that can be used for troubleshooting. Public-write access is needed because not all executables that write to this file run as root.

You can adjust the level of logging information provided in this log file by adjusting the vault logging level parameter on the **Logging** page of the master server's properties via **Host Properties** on the NetBackup Console.

Only root, or administrator, running on the master server, can execute this command.

EXAMPLES

Example 1:

To vault the profile `my_profile`, enter:

```
vltrun my_profile
```

Example 2:

The following command vaults the images for robot 0, Financials, and Weekly:

```
vltrun 0/Financials/Weekly
```

FILES

UNIX:

```
/usr/opensv/netbackup/vault  
/usr/opensv/netbackup/bp.conf  
/usr/opensv/netbackup/logs/bpbrmvt/log.mmddyy  
/usr/opensv/netbackup/logs/bpcd/log.mmddyy  
/usr/opensv/netbackup/logs/vault/log.mmddyy  
/usr/opensv/netbackup/db/vault/vault.xml  
/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx  
/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx/  
vltrun.log  
/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx/  
detail.log
```

Windows:

```
install_path\NetBackup\vault  
install_path\NetBackup\bp.conf  
install_path\NetBackup\logs\bpbrmvt\mmddyy.log  
install_path\NetBackup\logs\bpcd\mmddyy.log  
install_path\NetBackup\logs\vault\mmddyy.log  
install_path\NetBackup\db\vault\vault.xml  
install_path\NetBackup\vault\sessions\vault_name\sidxxx  
install_path\NetBackup\vault\sessions\vault_name\sidxxx\  
vltrun.log  
install_path\NetBackup\vault\sessions\vault_name\sidxxx\  
detail.log
```



EXIT STATUS

Vault may exit with a status code greater than 255. Such status codes are called extended exit status codes. For such a case, the exit status returned to the system is 252, and the actual exit status is written to `stderr` in the format, `EXIT status = exit status`

The extended exit status values are documented in the *NetBackup Troubleshooting Guide* and in the *NetBackup Troubleshooting Wizard*.

SEE ALSO

`vltadm.1m`, `vlteject.1`, `vltinject.1`, `vltoffsitemedia.1`, `vltopmenu.1`

Upgrading bpvault to NetBackup Vault 4.5

B

To upgrade from the bpvault product to NetBackup Vault 4.5, you must reconfigure your vault processes. If your Vault configuration is relatively simple, for example, if only one media server is specified for each dup_param file, you may choose to reconfigure on your own. However, we suggest you contact the VERITAS Enterprise Consulting Service (ECS) to help with the transition from bpvault 3.4 to NetBackup Vault 4.5 if any of the following are true:

- ◆ If your bpvault scripts were customized,
- ◆ If you were not involved with the initial bpvault setup,
- ◆ If your bpvault configuration has multiple media servers per dup_param file,
- ◆ If your bpvault configuration has multiple robots per dup_param file,
- ◆ If you are not comfortable with the guidelines presented in this appendix.

To contact ECS:

Tel: 1-650-527-8000 (outside US)

Tel: 1-800-327-2232 (inside US)

Fax: 1-650-527-8050

Introduction to Vault 4.5

Vault 4.5 is more tightly integrated into the NetBackup product. You can schedule Vault jobs through the NetBackup scheduler and view Vault jobs in the NetBackup Activity Monitor. A program named `vlt-run` has replaced the `bpvault.all` shell script.

Vault 4.5 contains most of the functionality of 3.4, but some of it has been implemented in a different way. Therefore, there is not a one-to-one correspondence between what you did for 3.4 and what you will do for 4.5. There are also some new features in 4.5. Refer to the tables below for details.

In Vault 4.5, the vaulting process is broken into three parts:

- ◆ Selection and duplication of images,



- ◆ Catalog backup,
- ◆ The eject and reporting process.

Vault 4.5 does not solve the resource contention and media contention problems that exist in bpvault. However, Vault 4.5 does contain new load-balancing logic which allows a vault session to utilize all drive pairs until the last of the backups is being duplicated. This feature should reduce contention issues.

If You Choose Not to Upgrade

bpvault 3.4 is compatible with NetBackup 4.5 as long as you do not increase the Maximum Backup Copies setting in NetBackup. The default setting is 2.

Feature Comparison between Versions

This table provides a synopsis of new or changed features

Feature	Version 3.4	Version 4.5
GUI	Not available	Available on Windows and UNIX platforms. See the NetBackup <i>Release Notes</i> for details.
Configuration of Vault processes through a GUI or a MUI	Not available	Available on Windows and UNIX platforms. See the NetBackup <i>Release Notes</i> for details.
Inline Tape Copy during Backup	Not available	Create up to 4 copies at a time
Inline Tape Copy during Duplication	Not available	Create up to 4 copies at a time
Number of backup copies possible	2	10
Number of retention levels possible	Maximum of 10	Maximum of 25
NetBackup Logging	Not available	Vault logs go directly to netbackup/logs

Feature	Version 3.4	Version 4.5
Disk staging	No automatic freeing of disk space	Create multiple tape copies from a disk image and optionally automatically free up disk space following successful duplication
Choose backup images based on client	No, only class (policy) and schedule	Yes, through the GUI or <code>vltadm</code>
Choose backup images based on backup type	No, only class (policy) and schedule	Yes, through the GUI or <code>vltadm</code>
Choose backup images based on media server	Yes	Yes, through the GUI or <code>vltadm</code>
Notify scripts	Yes, through VERITAS consulting.	Templates are provided in goodies directory.
Multi-eject support	Yes, through VERITAS consulting.	Available for robots with media access ports.
Consolidated eject for multiple Vault sessions	Not available	Available through CLI (<code>vlt eject</code>) or MUI (<code>vltopmenu</code>).
Activity Monitor Integration	Not integrated	Vault jobs and duplication jobs appear in Activity Monitor and can be cancelled.
Use of NetBackup Scheduler	Not available	Schedule Vault jobs through the NetBackup Scheduler
Cleanup of session files	Manual	Automatic
Vault all backups written by <i>all other</i> media servers.	<code>dup_cleanup</code> parameter	All media servers selected by default when choosing backups. Or, you can specify individual media servers.
Menu User Interface (MUI)	<code>bpvault.menu</code> (administration menu), <code>bpvault.opsmenu</code> (operations menu)	<code>bpvault.menu</code> has been retired <code>vltadm</code> (configuration menu) <code>vltopmenu</code> (operations menu)



Feature	Version 3.4	Version 4.5
Nutcracker and MKS Toolset required on NT systems	Required and installed with the product.	Not required
Configuration of Vault processes	dup_param files	Vault profiles configured via the Vault GUI or MUI
Retention level of duplicate copy	Calculated based on original retention level.	Not calculated. A duplicated copy can be given one of 25 retention levels.
Alternate read server to be used by the duplication process	server parameter ALTREADHOST option	GUI option labeled: Alternate read server. Read original backups using media servers that are different from the media server that wrote the backups.
Vaulting non-NetBackup media managed by Media Manager	Available through VERITAS consulting.	Use notify script, vlt_ejectlist_notify, which creates the addon_medialist file
Internationalization	Not available	Yes
Selection of backup images	Filtered by date range, class (policy), and schedule.	Filtered by date range, policy (class), schedule, client, media server, backup type, and source volume group
Preview	In bpvault.menu, run the Preview step.	vltrun <i>profile_name</i> -preview
Restart	Fix the offending problem, then manually step through the process using bpvault.menu	Fix the offending problem, then rerun the job. Vault figures out which images have already been duplicated.
Executables	bpvault.all	Compiled
Robots Supported	TL8, TLD, ACS	Must have a media access port.
Multi-tape catalog backup	Configure your catalog backup to use the two-step method.	Given the name of a backup policy, Vault 4.5 will execute the two-step method.

Feature	Version 3.4	Version 4.5
Alternate Media Server Names	server parameter althostname option	Vault Properties dialog

Feature Descriptions

This section describes some of the features in the table above in more detail.

Configuration File

`dup_param` files in `bpvault` have been replaced with profiles in Vault 4.5. For each `dup_param` file you have, you will need to create at least one Vault profile. You will configure each profile through the new Vault GUI or the `vltadm` MUI.

If you have a `dup_param` file which overrides more than one retention level, you must create a separate Vault profile for every original retention level override.

Refer to Chapter 5, “Configuring Vault” for more information.

Choosing Backup Images

`bpvault` filtered the images selected for duplication by class and schedule and wrote the results to `preview.list`.

In Vault 4.5, images are filtered by:

- ◆ Time Range
- ◆ Policy (Class)
- ◆ Schedule
- ◆ Client
- ◆ Media Server
- ◆ Type of Backup
- ◆ Source Volume Group

In both products, a backup image will be excluded if it has already been vaulted by the active profile.



Cleanup of Session Files

NetBackup 4.5 will automatically delete the session directories when they have aged a configured amount of time. (In 3.4 these would have been the DUPxxx directories). The default is 30 days. To change this, set the **Keep Vault Logs** parameter on the Global Attributes page of the master server's properties in **Host Properties** on the NetBackup console.

Cleanup Flag

In Vault 4.5, the original cleanup flag is effectively enabled whenever you are using the Advanced view of the Duplication tab. This is because the Advanced view requires that you specify a list of media servers. If the list does not include one of the media servers included in the Media Servers list on the Choose Backups tab, the behavior is the same as if you had set the cleanup flag in the `bpvault dup_param` file. To avoid cleanup, the source media servers listed in the advanced view of the Duplication tab must match the media servers specified on the Choose Backups tab.

Disk Staging

If your original backups reside on disk, Vault can copy them to one or more tapes. You can configure a profile to automatically expire the original disk backup image if, and only if, the duplication was successful. This action frees up disk space for future backups. This feature will not expire original backups which are on tape.

Retention Level Override

In bpvault 3.4 it is possible to apply a variance to the expiration date of images ejected, according to the original retention level of each image. Vault 4.5 provides additional (up to 25) configurable retention levels but does not use the retention level of the original backup to determine the retention level for the duplicated copy. There are two ways to handle this with NetBackup Vault 4.5.

With NetBackup 4.5 you can configure your backup policies to create more than one copy of each backup, where each copy would get a different retention level. Subsequently, your vault profile would vault one (or more) of those copies. If you have enough resources to operate in this fashion, we highly recommend that you do so.

Alternatively, you can configure Vault 4.5 to make a copy of your backup images, where the duplicated copy would be given a retention level that is different from the retention level of the original backup image(s). If you configure Vault 4.5 to make multiple copies of each backup image, then each of the (up to 4) copies can be assigned a different retention level.

Note If you have a `dup_param` file that adjusts more than one original retention level, to configure the exact behavior you will need to create a separate Vault 4.5 profile for each retention level used by the `dup_param` file. These separate profiles must further filter the backup images to reduce the preview list of backup images to those with the specific original retention level.

Alternate Media Server Names

In `bpvault`, the issue of multiple names for a media server was addressed in the `server` parameter. In Vault, you may specify a list of comma-separated alternate media server names. These names are case sensitive. To configure this via the GUI, use **Vault Properties** from the **Actions** menu on the NetBackup console.

Each of the names within any one entry are considered equal. If any one of the names in an entry is specified in a media server criterion on the Choose Backups tab, or as a source media server value on the Advanced Duplication tab, then Vault will search for all of the names in the server list when looking for the corresponding criteria.

Vaulting non-NetBackup Media Managed by Media Manager

You may specify additional media IDs to be ejected by creating a file named `addon_medialist` at a specific point in the Vault process. These media IDs are not case sensitive. In `bpvault` 3.4 you had to modify the controlling shell script to add media IDs. In Vault 4.5, a notify script, `vlt_ejectlist_notify`, is provided to perform this function. Vault will append the media specified in the `addon_media` list to the total list of media to be ejected. Vault also tracks this non-NetBackup media.

Multiple Robots

If you have multiple robots attached to your media server, and you want to configure a profile which makes duplicates of backups which are on media in only one of the robots, you must specify the robot. This situation was handled by the `multrobots_onserver` parameter in `bpvault` 3.4. In Vault 4.5, you set the **Source volume group** field on the Choose backups tab.

Preview

Using `vltrun profile_name -preview` gives you the opportunity to sanity-check your configuration on the **Choose Backups** tab. This parameter starts a new vault job, performs a search on the image catalog based on the criteria specified on the **Choose Backups** tab, and then exits. No actions are performed on the backup images selected.



After running this command, you will need to access the following file to view the results, that is, the images selected based on the criteria specified.

```
UNIX: /netbackup/vault/sessions/vault_name/sidxxx/preview.list

Windows:
install_path\NetBackup\vault\sessions\vault_name\sidxxx\
preview.list
```

For each backup image selected, this file contains information including the time of the backup, the backup policy name, the schedule name, the backup ID, and the media server.

Note This list of images may be a superset of the list of backup images that would be vaulted.

Depending on how your profile is configured, two more criteria may be applied to this list before the actual processing of the images to be vaulted occurs.

- ◆ If the Duplication step is configured to only duplicate disk images, any backups which appear in the list but do not have a disk copy, but only have a copy on removable media, such as tape, will not be vaulted.
- ◆ When the Eject step begins for this profile, any backups in the list which do not have a copy on media in one of the off-site volume pools listed for the Eject step will not be vaulted.

Where to Find Data Between Versions

The tables below provide guidance in understanding the differences and similarities between bpvault 3.4 and Vault 4.5.

Comparison of Variables and Actions in Each Version

This table will help you move the information in a bpvault 3.4 parameter file to a Vault 4.5 format using the new Vault GUI.

Robot Identifier Properties	
Vault 4.5	bpvault 3.4
Robot number	robot_num
Vault Identifier Properties	
Vault name	vault
Vault vendor	vault_vendor
Robotic volume group	robot_group
Off-site volume group	vault_group
First off-site slot ID	starting_slot

Profile Properties	
Vault 4.5	bpvault 3.4
Choose Backups Tab	
Backups Started Between X days Y hours ago	<i>duplicate_days</i> (X) <i>duplicate_hours</i> (Y)
And X days Y hours ago	<i>days_startfrom</i> (X) <i>hours_startfrom</i> (Y)
Type of backup	N/A
Source volume group	<i>multrobots_onserver_srcgrp</i>
Clients	N/A
Media servers	N/A
Backup policies	<i>class</i>
Schedules	<i>schedule</i>
Duplication Tab	
Source backups reside on	N/A
Number of read drives	<i>server, dstunit</i>
Read original backups...	Server listed as ALTREADHOST on <i>server</i> parameter
Multiple copies	N/A
Storage unit	<i>dstunit</i>
Write drives	Number following storage unit in <i>dstunit</i>
Volume pool	<i>pool</i>
Retention level	<i>retention_length</i>
Make this copy primary	N/A
Preserve multiplexing	<i>mpxdup</i>
Expire original disk backup...	N/A
Catalog Backup Tab	
Media server	<i>dbbackuphost</i>
Volume pool	<i>dbpool</i>
Retention period	<i>days_holddbtape</i>
Number of (serial) Catalog...	<i>num_dbdups</i>
Files to be backed up	<i>dbpath</i>
Backup policy for multiple-tape...	N/A
Eject Tab	
Eject tapes from these volume pools only	<i>pool</i>
Suspend media on which backups were written between X and Y days ago	<i>suspendmedia</i> (checkbox) <i>days_suspend</i> (X)
Eject mode	<i>robot_eject</i>
Reports Tab	
Report header	<i>eject_header</i>
Reports for media going off site Reports for media coming on site Detailed media reports	Manual editing of bpvault.reports
Recovery Report for Vault between X and Y days ago	Manual editing of bpvault.reports (checkbox) <i>recoverlength_days</i> (X) <i>recoveryto_days</i> (Y)



Change Report Titles	<i>compinv_header</i> <i>distdtl_header</i> <i>distlib_header</i> <i>distsum_header</i> <i>distvlt_header</i> <i>fullvlt_header</i> <i>involt_header</i> <i>origdtl_header</i> <i>origejc_header</i> <i>picklib_header</i> <i>pickvlt_header</i> <i>recovery_header</i>
E-mail	<i>run_report_mail</i> (checkbox) MAILNAMES environment variable in bpvault.env (text box)
Print command	<i>run_report_print</i> (checkbox) LPR environment variable in bpvault.env (text box)
Folder	<i>run_report_file</i> (checkbox) <i>report_dir</i> (text box)
Report mode	<i>run_report</i>
Actions Menu	
Vault 4.5	bpvault 3.4
E-mail address for notification of session status	N/A
Default e-mail address for reports	MAILNAMES environment variable in bpvault.env
Alternate Media Server Names	Listed after server name and number of drive pairs in <i>server</i> parameter
Outside GUI	
Vault 4.5	bpvault 3.4
VAULT_VERBOSE in bp.conf	<i>debug</i>

Comparison of Parameters in Each Version

This table contains bpvault 3.4 parameters and their equivalents in Vault 4.5.

Parameter in bpvault 3.4	Location in Vault 4.5
<i>acs</i>	Automatically determined by GUI
<i>acs_ack_timeout</i>	No longer used
<i>acs_cap_0</i>	Automatically determined by GUI

<i>acs_cap_1</i>	Automatically determined by GUI
<i>acs_cap_count</i>	Automatically determined by GUI
<i>acs_cap_num</i>	Automatically determined by GUI
<i>acs_command_timeout</i>	No longer used
<i>acs_csi_hostname</i>	Automatically determined by GUI
<i>acs_lockfile_prefix</i>	No longer used
<i>acs_lsm</i>	Automatically determined by GUI
<i>acs_nodrives_onmaster</i>	No longer used
<i>acs_number_tries</i>	No longer used
<i>acs_sockname</i>	No longer used
<i>assignslot_multivault</i>	No longer used; enabled by default
<i>backupdb_command</i>	No longer used
<i>bpexpdate_command</i>	No longer used
<i>bpmedia_command</i>	No longer used
<i>bpmedia_option</i>	No longer used; only suspend option supported
<i>bpvault_dir</i>	Directory is always .../netbackup/vault/sessions/ <i>vault_name</i>
<i>change_command</i>	No longer used
<i>changeexpdate</i>	No longer used; set retention levels of duplicates via Duplication tab
<i>class_sort_field</i>	No longer used
<i>class</i>	See Choose Backups Tab
<i>classexclude</i>	No longer used; set exclude policies via Choose Backups
<i>cclist_command</i>	No longer used
<i>command_log</i>	No longer used
<i>compinv_header</i>	See Reports Tab
<i>copy_number</i>	No longer used
<i>dasadmin_client</i>	Automatically determined by GUI
<i>dasadmin_command</i>	No longer used
<i>dasadmin_eject</i>	Automatically determined by GUI
<i>dasadmin_eject_comple te</i>	Automatically determined by GUI
<i>dasadmin_insert</i>	Automatically determined by GUI
<i>dasadmin_server</i>	Automatically determined by GUI
<i>days_holddbtape</i>	See Catalog Backup Tab
<i>days_startfrom</i>	See Choose Backups Tab
<i>days_suspend</i>	See Eject Tab
<i>dbbackuphost</i>	See Catalog Backup Tab
<i>dbpath</i>	See Catalog Backup Tab
<i>dbpool</i>	See Catalog Backup Tab
<i>debug</i>	See Outside GUI above
<i>debug_counter</i>	No longer used
<i>distdtl_header</i>	See Reports Tab
<i>distlib_header</i>	See Reports Tab
<i>distsum_header</i>	See Reports Tab
<i>distvlt_header</i>	See Reports Tab
<i>drive_pairs</i>	No longer used
<i>dstunit</i>	See Duplication Tab
<i>dup_cleanup</i>	No longer used; enabled by default, use Media server list in Choose Backups to only allow specific servers to be duplicated
<i>dup_file_template</i>	No longer used



<i>duplicate_command</i>	No longer used
<i>duplicate_days</i>	See Choose Backups Tab
<i>duplicate_hours</i>	See Choose Backups Tab
<i>eject_command</i>	No longer used
<i>eject_command_log</i>	No longer used
<i>eject_header</i>	See Reports Tab
<i>eject_repeat</i>	No longer used
<i>error_file_template</i>	No longer used
<i>fullvolt_header</i>	See Reports Tab
<i>hours_startfrom</i>	See Choose Backups Tab
<i>imlist_command</i>	No longer used
<i>involt_header</i>	See Reports Tab
<i>line_max</i>	No longer used
<i>log_file_template</i>	No longer used
<i>media_type</i>	Automatically determined by GUI
<i>medlist_command</i>	No longer used
<i>mmhost</i>	Automatically determined by GUI
<i>mpxdup</i>	See Duplication Tab
<i>mtlib_bulk</i>	Automatically determined by GUI
<i>mtlib_command</i>	No longer used
<i>mtlib_device</i>	Automatically determined by GUI
<i>mtlib_host</i>	Automatically determined by GUI
<i>multrobots_onserver</i>	No longer used; enabled by default
<i>multrobots_onserver_sr_cgrp</i>	See Choose Backups Tab
<i>nb_version</i>	No longer used
<i>no_message_minutes</i>	No longer used
<i>num_dbdups</i>	See Catalog Backup Tab
<i>origdtl_header</i>	See Reports Tab
<i>origejc_header</i>	See Reports Tab
<i>picklib_header</i>	See Reports Tab
<i>pickvlt_header</i>	See Reports Tab
<i>pool</i>	See Duplication Tab
<i>print_log_flag</i>	No longer used
<i>query_command</i>	No longer used
<i>recoverlength_days</i>	See Reports Tab
<i>recovertv_days</i>	See Reports Tab
<i>recovery_header</i>	See Reports Tab
<i>remote_command</i>	No longer used
<i>report_dir</i>	See Reports Tab
<i>retention_length</i>	No longer used; set retention levels of duplicates via Duplication tab
<i>robot_eject</i>	See Eject Tab
<i>robot_group</i>	See Vault Identifier Properties
<i>robot_inventory</i>	No longer used
<i>robot_num</i>	See Robot Identifier Properties
<i>robot_type</i>	Automatically determined by GUI
<i>run_report</i>	See Reports Tab
<i>run_report_file</i>	See Reports Tab
<i>run_report_mail</i>	See Reports Tab

<i>run_report_print</i>	See Reports Tab
<i>schedule</i>	See Choose Backups Tab
<i>server</i>	See Duplication Tab
<i>sleep</i>	No longer used
<i>starting_slot</i>	See Vault Identifier Properties
<i>suspendmedia</i>	See Eject Tab
<i>tl8_count</i>	Automatically determined by GUI
<i>tl8_eject</i>	No longer used
<i>tl8test_bus</i>	Automatically determined by GUI
<i>tl8test_command</i>	No longer used
<i>tl8test_host</i>	Automatically determined by GUI
<i>tl8test_hun</i>	Automatically determined by GUI
<i>tl8test_port</i>	Automatically determined by GUI
<i>tl8test_robotpath</i>	Automatically determined by GUI
<i>tl8test_target</i>	Automatically determined by GUI
<i>tld_count</i>	Automatically determined by GUI
<i>tlh_count</i>	Automatically determined by GUI
<i>tlm_count</i>	Automatically determined by GUI
<i>tpreq_command</i>	No longer used
<i>tpunmount_command</i>	No longer used
<i>vault</i>	See Vault Identifier Properties
<i>vault_group</i>	See Vault Identifier Properties
<i>vault_type</i>	No longer used; select Skip Duplication on Duplication tab to vault originals
<i>vault_vendor</i>	See Reports Tab
<i>vmchange_timeout</i>	No longer used
<i>volmgr_miscdir</i>	No longer used

Cross References: Executables, Directories, and Files

This table provides a quick reference for renamed executables, directories, and files.

Version 3.4	Version 4.5
bpvault.all	vltrun
bpvault.opsmenu	vltopmenu
changecopy.sh	bpchangeprimary
bpvault.report	vlteject
*.eject	vlteject



Version 3.4	Version 4.5
netbackup/vault/production	netbackup/bin
netbackup/vault/bin	N/A
netbackup/vault/vaultname/DUPxxx	netbackup/vault/sessions/vaultname/sidxxx



Recovering Damaged Media

C

This appendix provides information on how to use duplicated images to replace damaged media.

Note This image recovery process assumes that the NetBackup system and image catalog are current and up-to-date.

Steps in Recovering Backup Images

The most frequent reason to use a duplicated tape is that the original backup tape for a specific image or set of images is lost or damaged. The steps below represent an outline of the recovery procedure. Each step is then discussed in detail.

1. Identify the damaged tape.
2. Determine which images were originally on the damaged tape.
3. Determine which duplicate tapes contain these images.
4. Tell NetBackup to use the duplicate tape to restore these images.
5. Freeze the tapes containing the duplicated versions so they won't expire.
6. Request that the tapes be returned from the off-site vendor.
7. Place the tapes into the robot, and notify Media Manager of their new location.
8. Perform a normal restore, which will automatically request the mount of the duplicated versions.
9. Unfreeze the duplicated tapes to allow the image/media to expire normally.
10. Optionally, make new duplicates to replace damaged ones.



Following the Image Recovery Procedure

Most of the commands listed below can be found in the following directory:

UNIX: `usr/opensv/netbackup/bin/admincmd`

Windows: `install_path\netbackup\bin\admincmd`

We recommend that you place this directory in your search path when you are engaged in image recovery.

Identifying the Damaged Media

When you receive an error message during a restore, the errors are logged to the restore log and also show up on the Activity Monitor as the restore fails. You can set up a procedure using NetBackup scripts to send errors to an event management console to notify the storage administrator immediately of this type of media error.

Determining Which Backup Images Were on the Damaged Tape

All images on a specific tape can be identified by running the `bpimmedia` command. This command will scan the whole NetBackup image catalog, so it may take a few minutes depending on the size of that catalog. For example:

UNIX: `bpimmedia -mediaid TAPEID`

Windows: `install_path\NetBackup\bin\admincmd\bpimmedia.exe
-mediaid TAPEID`

Here's an example of the output. It shows that there is one image on the tape "S05423" for the client "fgolddust". It also shows that this image has been duplicated since it has (FRAG 2) entries. The full image name is "fgolddust_0862806643":

UNIX: `nirvana# bpimmedia -mediaid S05423`

Windows: `C:\Veritas\NetBackup\bin\admincmd bpimmedia.exe -mediaid
S05423`

```
IMAGE fgolddust 2 fgolddust_0862806643 goldust_BR1 0 Full_Weekly 0
3 19360 8654 85043 0 0
FRAG 1 -1 2293 0 2 6 2 S05423 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 1 232848 0 2 6 1 S02643 nirvana 64512 2 862804830 3 0 *NULL*
FRAG 1 2 1225539 0 2 6 2 S02643 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 3 70182 0 2 6 3 S02643 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 4 825700 0 2 6 1 S05423 nirvana 64512 2 862808446 3 0 *NULL*
FRAG 2 -1 2293 0 2 6 2 S04440 nirvana 32768 0 862927577 2 0 *NULL*
```

```
FRAG 2 1 2335584 0 2 6 1 S04440 nirvana 32768 2 862927577 2 0
*NULL*
```

Determining Which Duplicate Tapes Were Used and Their Host

In step (2) above, the (FRAG 2) entries show that an image has been duplicated. The (FRAG 2 1) entry is the duplicate copy, whereas on copy 1, there were 4 fragments (usually due to multiplexing). The (FRAG 2 -1) entry is the TIR duplicate. In this case, the image fgolddust_0862806643 is using media S04440 for duplicating all of the original fragments. This is normal because the original image was multiplexed onto 4 tapes, while the duplicate was de-multiplexed during image duplication, and could fit on one tape.

Also note that the host for the media is printed for each fragment, in this case nirvana. With media servers, the host could be different than the master. Under Vault, the duplication should normally occur on the same server that made the original backup, so the host server names would be the same for both copies of the image.

You can confirm this information by using bpimagelist:

```
UNIX: nirvana# bpimagelist -backupid fgolddust_0862806643
Windows: C:\Veritas\NetBackup\bin\admincmd bpimagelist.exe
-backupid fgolddust_0862806643

IMAGE fgolddust 0 0 2 fgolddust_0862806643 golddust_BR1 0 *NULL*
root Full_Weekl
y 0 3 862806643 4591 865485043 0 0 2356562 19360 2 7 1
golddust_BR1_0862806643_F
ULL.f *NULL* *NULL* 0 1 0 2 865830643 *NULL* 1 0 0 0 0 *NULL*
HISTO -1 -1 -1 -1 -1 -1 -1 -1 -1 -1
FRAG 1 -1 2293 0 2 6 2 S05423 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 1 232848 0 2 6 1 S02643 nirvana 64512 2 862804830 3 0 *NULL*
FRAG 1 2 1225539 0 2 6 2 S02643 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 3 70182 0 2 6 3 S02643 nirvana 64512 0 862804830 3 0 *NULL*
FRAG 1 4 825700 0 2 6 1 S05423 nirvana 64512 2 862808446 3 0 *NULL*
FRAG 2 -1 2293 0 2 6 2 S04440 nirvana 32768 0 862927577 2 0 *NULL*
FRAG 2 1 2335584 0 2 6 1 S04440 nirvana 32768 2 862927577 2 0
*NULL*
```

To confirm which is the primary copy (the copy to be used for restores), use the -L option with bpimagelist:

```
UNIX: nirvana# bpimagelist -L -backupid fgolddust_0862806643 | grep
Primary
Primary Copy: 1
```



```
Windows: C:\Veritas\NetBackup\bin\admincmd bpimagelist.exe -L
-backupid fgolddust_0862806643 | grep Primary
Primary Copy: 1
```

Telling NetBackup to Use the Duplicated Copy Rather Than the Original

The command `bpimage -npc` is executed to select which image is used for restoring an image:

```
UNIX: nirvana# bpimage -npc 2 -backupid fgolddust_0862806643
-client fgolddust
Windows: C:\Veritas\NetBackup\bin\admincmd bpimage.exe -npc 2
-backupid fgolddust_0862806643 | client fgolddust
```

To confirm the new primary copy:

```
UNIX: nirvana# bpimagelist -L -backupid fgolddust_0862806643 | grep
Primary
Primary Copy: 2
Windows: C:\Veritas\NetBackup\bin\admincmd bpimage.exe -npc 2
-backupid fgolddust_0862806643 -client fgolddust
```

Freezing the Duplicated Copy to Ensure Restore

Use the command `bpmedia -freeze` to keep NetBackup from expiring the images on the media and to keep the media assigned in Media Manager. You should also use the media host for this image, which was printed by `bpimage` in step (2) above. This is required when the host is different than the machine on which you are running this command.

```
UNIX: nirvana# bpmedia -freeze -ev S04440 -host nirvana
Windows: C:\Veritas\NetBackup\bin\admincmd bpmedia.exe -freeze -ev
S04440 -host nirvana
```

If a restore was initiated at this point, media ID S04440 would be used for the restore.

Requesting the Media to be Returned from the Vault

Document the process for requesting media from your off-site vendor keep it with information for the operations team. In most cases, you will need to fill out a form or contact the vendor by phone to request that a particular piece of media to be returned. You must specify the media ID of the tapes to be used for restores. The off-site vendor may

also require an account number, vault ID, and the slot in which the media physically resides. As seen below, you may find most of this information in Media Manager, by using the command `vmquery`, located in the following directory:

UNIX: `/usr/opensv/volmgr/bin`

Windows: `install_path\volmgr\bin`.

For example, the slot number can be found by querying the Media Manager and looking at the description field. In this case it is S278:

```
UNIX: vmquery -m S04440
Windows: C:\Veritas\Volmgr\bin vmquery.exe -m S04440
=====
media ID:S04440
media type:8MM cartridge tape (4)
barcode:-----
description:CH_V1|101|S278|00000000
volume pool:Duplicates (3)
robot type:NONE - Not Robotic (0)
volume group:vault_grp
created:Tue Sep 3 10:08:32 2000
assigned:Tue May 6 00:11:45 2001
last mounted:Tue May 6 11:34:25 2001
first mount:Tue Sep 3 18:20:48 2000
expiration date:---
number of mounts:21
max mounts allowed:---
status:0x0
=====
```

Placing Returned Tape(s) Back into the Robot

Once the tape is returned from the off-site vendor, you must place it in the appropriate robotic library. Consult the operations team for details on how to place tapes into the library.

After the media has been placed into the MAP, from the NetBackup Administration Console, choose **Media and Devices Management**. Choose the **Inventory Robot...** option. Select the **Empty Media Access Port Prior to Update** checkbox.

You can also perform this function using the `vlthinject` command.

Performing Normal Restore

At this point, a normal restore of the same image should use the duplicate media rather than the original media. The restore log should show a mount request for the duplicate media.



Unfreezing Media Used for Duplicates

Once the restore is successful, unfreeze the duplicate media to allow the normal expiration process to be followed. If you want to send the tape off-site again, either remove it from robot or leave it in the robot as the primary copy. We recommend you suspend the media so that future images do not get written to it.

```
UNIX: nirvana# bpmedia -unfreeze -ev S04440 -host nirvana
Windows: C:\Veritas\NetBackup\bin\admincmd bpmedia.exe -unfreeze
        -ev S04440 -host nirvana
```

Creating New Duplicate Images

See “Bad or Missing Duplicate Tape” on page 136 in the Troubleshooting chapter.

Modifying the NetBackup Catalog for a Large Number of Images

In a disaster recovery situation where a large number of images need to have their primary copy modified, run the `bpchangeprimary` command. This command will change the primary copy of all the backup images in the off-site volume pool for which the media was returned from the off-site vault.

Revaulting Media After a Restore

When you bring media back from an off-site vault to do a restore, revaulting those tapes is a manual procedure:

1. We suggest that you write-protect the tapes, either by physically write-protecting the media (for instance, by flipping the write-protection switch) or by using the `-freeze` option.
2. Use `vltinject` to replace the media in the robot.
3. Make sure that the copy is primary.

If you need to restore because your on-site tape was damaged or lost, use the catalog node on the NetBackup console to make another copy of your data to keep on site. Then use `bpchangeprimary` to mark the media primary. For more information about `bpchangeprimary`, see the NetBackup *System Administrator's Guide*.

4. Restore your data.
5. Manually eject the media, using one of the following methods:

- Use the media manager command line interface `vmchange` command.
- Highlight the media ID for the media you are working with. Then select the **Eject Volumes from Robot....** operation on the **Actions** menu in the NetBackup Administration Console.

Note `vlteject` and `vltopmenu` will not work for this purpose.

6. Manually transfer the media to the off-site volume group, using one of the following methods:
 - Use the media manager command line interface `vmchange` command.
 - Highlight the media ID for the media you are working with. Then select the **Change Volume Group....** operation on the **Actions** menu in the NetBackup Administration Console.
7. Remove write-protection from the media.
8. Return the media to your vault vendor so that all backups on that media will be available for future disaster recovery.
9. Run the *Recovery* report to ensure that the media is available for future disaster recovery operations.





Vault’s File and Directory Structure

Vault is installed in the directory specified by the *install_path*, on UNIX */usr/opensv/netbackup*.

Directories and Files Created During Installation

The following directories are created during the installation:

DIRECTORY NAME	UNIX PERM	UNIX GROUP
help/vltadm	0755	bin
db	0755	bin
db/vault	0755	bin
vault/sessions	0777	bin

The table below shows the subdirectories and files in each of these directories. Some are created when Vault is installed, others are created as Vault sessions run.

Vault Programs and Scripts	Purpose
<i>install_path</i> \netbackup\vault\	Contains programs, working directories, etc.
<i>install_path</i> \netbackup\vault\sessions	A subdirectory containing working session directories and log files.
<i>install_path</i> \netbackup\vault\sessions\sidxxx	Subdirectory containing working session subdirectories. Can be manually removed to reduce disk usage if necessary.



../bin/bpbrmvlt	Process that kicks off vltrun from a scheduled or manual vault policy.
../bin/goodies/vlt_ejectlist_notify	Script called by the vault session just before vault tapes are ejected.
../bin/goodies/vlt_end_notify	Script called by the vault session just before it exits.
../bin/goodies/vlt_endeject_notify	Script called by the vault session at the end of eject processing.
../bin/goodies/vlt_start_notify	Script called by the vault session after it starts.
../bin/vltadm	Utility which has a menu interface that an administrator can use to configure NetBackup Vault and monitor its operations. vltadm requires root (administrator) privileges.
../bin/vltcore	Process that executes all the NetBackup commands. Executed multiple times during a session and called by vltrun.
../bin/vlteject	Command used to eject media from Vault sessions and run the reports selected in the profile.
../bin/vltinject	Command used inject media into a robot and update the Media Manager database.
../bin/vltoffsitemedia	Command which allows the user to change the off-site parameters of a given piece of media.
../bin/vltopmenu	Utility which allows the user to invoke a menu screen containing the various options that an operator of the NetBackup Vault feature can use. It allows the user to eject or inject media, print various reports individually or collectively, and consolidate reports and ejections across sessions.
../bin/vltrun	Process which drives a NetBackup Vault session by issuing a sequence of calls to the Vault engine (vltcore).
../help/vltadm/Catalog_Backups	Help file for vltadm.
../help/vltadm/Choosing_Backups	Help file for vltadm.
../help/vltadm/Duplication	Help file for vltadm.
../help/vltadm/Eject	Help file for vltadm.

../help/vltadm/Help	Help file for vltadm.
../help/vltadm/Main	Help file for vltadm.
../help/vltadm/Preferences	Help file for vltadm.
../help/vltadm/Reports	Help file for vltadm.
../help/vltadm/Robots	Help file for vltadm.
../help/vltadm/Tutorial	Help file for vltadm.
../help/vltadm/Vaults	Help file for vltadm.
..\sessions\vault_name\session.last	Counter to show current duplication session
..\sessions\vault_name\sidxxx\preview.list	Created by Vault as a first step in duplication.
..\sessions\vault_name\sidxxx\vault.error.file	Error log for other administrative commands performed by Vault; this file should be checked in case of problems.
..\sessions\vault_name\sidxxx\vault.err_suspend	Error log for other administrative commands performed by bpvault during suspend mode; this file should be checked in case of problems.
..\sessions\vault_name\sidxxx\vault.err	Error log for duplication of specific image.
..\sessions\vault_name\sidxxx\duplicate.log.nn	Contains output from bpduplicate.
..\sessions\vault_name\sidxxx\class.inventory	Listing of all configured NetBackup classes; for use with recovery report.
..\sessions\vault_name\sidxxx\preview.list	Duplication preview output file. Shows images which are to be duplicated.
..\sessions\vault_name\sidxxx\preview.list_suspend	List of images for which media will be suspended.
..\sessions\vault_name\sidxxx\image.list	NetBackup image catalog information for each image duplicated.
..\sessions\vault_name\sidxxx\image.list_suspend	NetBackup image catalog information for each image whose media will be suspended.



..\sessions\vault_name\sidxxx\media.list	NetBackup media used for originals and duplicates.
..\sessions\vault_name\sidxxx\media.list_suspend	NetBackup media used for originals to be suspended.
..\sessions\vault_name\sidxxx\nb_media.list	Contains the number of images, size of images, and expiration dates for original and duplicated media.
..\sessions\vault_name\sidxxx\nb_media.list_suspend	Contains the number of images, size of images, and expiration dates on original media to be suspended. Used for reports.
..\sessions\vault_name\sidxxx\nbudb.media	Contains media IDs that were allocated for NetBackup database backups.
..\sessions\vault_name\sidxxx\rcvrimage.inventory	NetBackup image catalog information for all classes between dates specified in the profile. For use with the recovery report.
..\sessions\vault_name\sidxxx\robot.inventory	Lists all the media currently residing in the robot (one media ID per line).
..\sessions\vault_name\sidxxx\returned_media.list	Temporary file used for expiring recalled tapes from offsite vault.
..\sessions\vault_name\sidxxx\volume.list	Media Manager inventory for duplicate pool and NetBackup database duplicate pool.
..\sessions\vault_name\sidxxx\volume.inventory_suspend	Media Manager inventory for original media pools.
..\sessions\vault_name\sidxxx\volume.db.list	Media Manager inventory for NetBackup database duplicate pool.
..\sessions\vault_name\sidxxx\volume_full.list	Media Manager inventory of all media.
..\sessions\vault_name\sidxxx\detail.log	Shows the output of every command that was executed during the session.
..\sessions\vault_name\sidxxx\duped.images	List of images successfully duplicated during the session.
..\sessions\vault_name\sidxxx\eject.list	List of media to be ejected for the session.

..\sessions\vault_name\sidxxx\reportout.nnn	Output of the reports generated during the session.
..\sessions\vault_name\sidxxx\summary.log	Concise view of detail.log listing major events during the session, such as how many images were copied. This log is appended for email notification.
..\sessions\vault_name\sidxxx\vltrun.log	Contains the commands executed in the session.





This functional design document provides an architectural services-oriented approach to building and maintaining a client/server based vault management system. This document defines the functional requirements, develops an architecturally-integrated set of services to meet those requirements, and documents the functional capacity for each service. To provide additional technical detail, this functional design further defines the technical components and technical design needed to provide these services. A final section lists operational procedures needed to deliver each service and assigns basic levels of staff responsibility.

Functional Design Overview

Overview

The chief benefit of this functional design is the consistent information for business, technical, and operational viewpoints. High level architectural service design provides a business understanding and real-world value. Technical implementation considerations are shown in the technical component diagrams and definitions. Procedural charts provides a hands-on understanding for operational staff and clarifies areas of responsibility.

- ◆ *Architectural Services:* A set of interrelated services is designed to meet the overall goals and/or challenges. This service model provides a framework for all functional features, and creates a visual model for further technical discussion. The services section uses tables and diagrams to create an organized and expandable system for future enhancements.
- ◆ *Technical Components:* This -section ties together the conceptual services with the actual NetBackup software and/or other software as needed. These diagrams and tables show how services are implemented and fundamentally interrelate. Tools definitions provide a technical design summarization for architectural considerations, and includes basic interface specifications.



- ◆ *Operational Procedures:* This final sub-section shows how the tools are to be used. The procedural diagrams and tables provide a real-world understanding, externalizing any assumptions about who would use the tools and in what manner.

Other Related Services

All storage management services are interdependent in some degree or another. Vault management relies upon both NetBackup and Media Manager services. It is designed to integrate with other operational services subsystems, such as Event Management and Help Desk. The interdependency on these other operational services creates the functional need for close integration of storage management services. Different functional designs provide detailed information for the different storage management sub-systems. It is beyond the scope of this functional design to document these other service areas, such as Help Desk. However, their value to an production environment is noted to provide an operational context.

Other Related Vault Documents

The NetBackup *Vault 4.5 System Administrator's Guide* provides basic installation, configuration and troubleshooting information. The NetBackup *Vault 4.5 Operator's Guide* provides day to day procedures to follow to work with vault reports, tapes and vault vendors.

Architectural Services

A vault management system is a server-centric approach that provides both backup duplication and off-site storage and retrieval of media. A Vault management system duplicates backup images onto tape or other media and simplifies restoring the duplicated files when the original backup image media is damaged or unavailable. A master-client implementation extends the storage to other machines by centrally controlling duplication for multiple backup servers simultaneously. A short list of basic services includes:

- ◆ Additional backup protection by duplicating backup images.
- ◆ Optional vaulting of original images/tapes.
- ◆ Direct support for different media types for both backup and duplication media.
- ◆ Backup of images onto different media types to support optimal cost-effective configurations.
- ◆ Maintenance of file OS level information and security.
- ◆ Scalable, distributed and heterogeneous implementation.

- ◆ Control of tape storage within tape robot vs. tapes ejected from robot for transference to off-site vault.
- ◆ Assignment of vault slot id location when required for use by vault vendor.
- ◆ Appropriate reports
- ◆ Use of existing NetBackup capabilities for key features to ensure compatibility and robustness:
 - Duplication image catalog for Recovery services.
 - Use of existing Media Manager services for fundamental media and robotic management and control.
 - Use of existing Media Manager database for keeping track of media containing duplicated images.
 - Use of NetBackup backup image catalog to keep track of which images need to be duplicated and which images are already duplicated.
 - Use of NetBackup media catalog to determine expiration data of used media.

The overall architectural design is defined as a distributed client/server system.

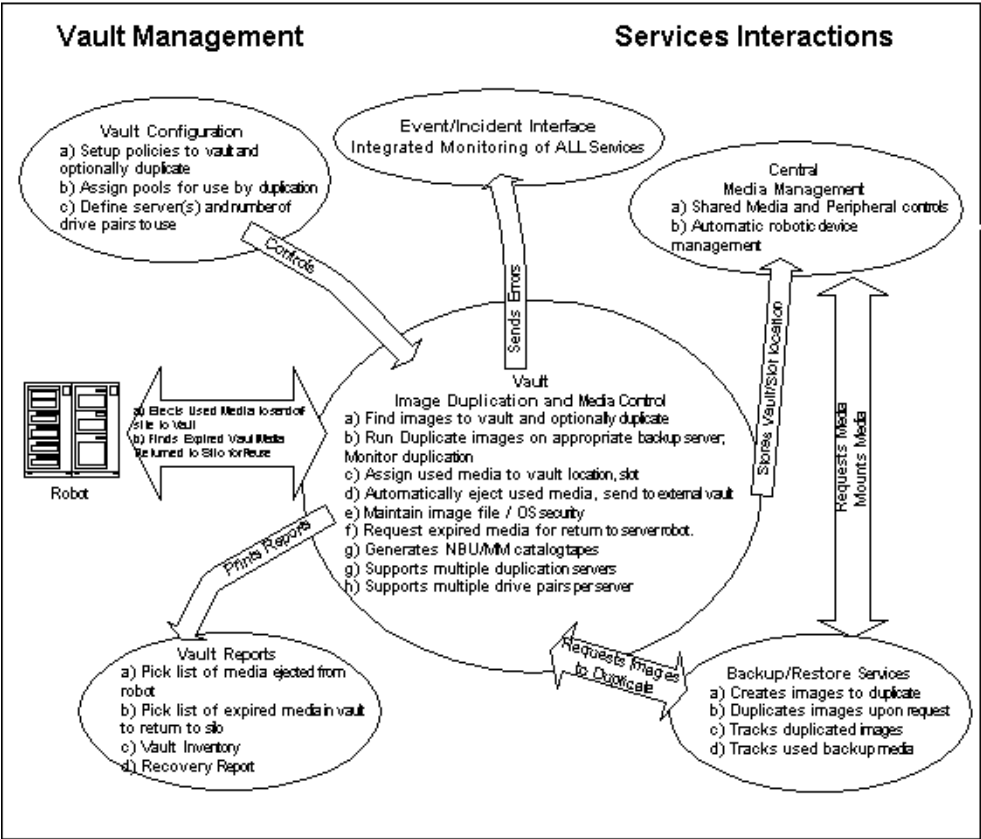
Client/server systems provide several basic features:

- ◆ Built-in network support. All services must exist within a network topology and thus do not require special configuration or software such as found in older, standalone designs.
- ◆ Scalable support for large numbers of backup servers. Server based features support faster processors and faster devices.
- ◆ Peer to peer controls. Servers can control other servers to provide better load balancing and different network topologies and bandwidths. Distributed servers ensure better production support and redundancy.
- ◆ Remote installation, control and configuration. Centralized management reduces management costs by making it easier to setup and run basic backup operations.



Services Interactions Diagram

The following diagram provides an architectural overview of the basic client/server backup design:



Client/Server Architectural Services

The overall set of client/server architectural services is provided in this chart:

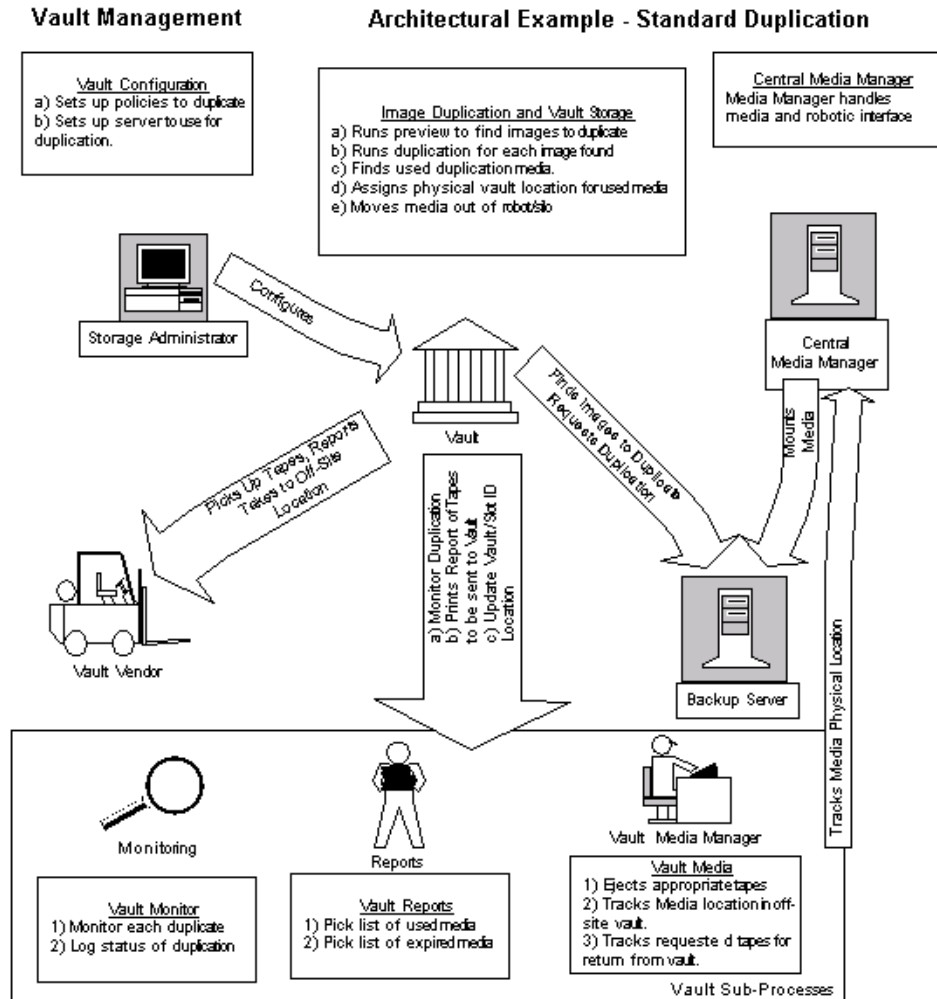
Hierarchical Storage Management		Architectural Services
Service	Business Challenge	Functional Capability
Vault Duplication	Protect Information - make sure appropriate images are duplicated Enterprise Scalability - support multiple servers	<ul style="list-style-type: none"> ◆ Determine backup images to duplicate. ◆ Duplicate images on multiple drives, multiple servers. ◆ Duplicate during day without using production network bandwidth. ◆ Optional vaulting of original images without duplication requirement. ◆ Can duplication locally, across LAN or WAN.
Vault Monitoring	Fast response	<ul style="list-style-type: none"> ◆ Monitor duplication process for successful completion ◆ Support interface to Event Management
Vault Configuration		Setup use of servers, which images to duplicate, other options.
Vault Reports	Protect Information - make sure media location is known	<ul style="list-style-type: none"> ◆ Print appropriate reports for sending media off-site to vault and returned from vault. ◆ Print regular inventory of vault.
Vault Media Management	Reduce Costs - reduce manual administration of media	<ul style="list-style-type: none"> ◆ Keep track of used media location in vault. ◆ Eject used media from robot for duplication session. ◆ Find returned, expired media and re-use in robot



BackupRestore Services	Protect Information	<ul style="list-style-type: none"> ◆ Creates duplicate on image by image basic. ◆ De-multiplexes backup image during duplication. ◆ Keeps track of both backup images and duplicated copy ◆ Keeps track of media used by duplicated copy. ◆ Simple image catalog change to restore from copy.
Media Manager Interface	Protect Information	<ul style="list-style-type: none"> ◆ Maintains media used information ◆ Maintains vault location information.

Architectural Example

This chart shows an example of how these services actually work. Detailed technical examples are provided in the technical components section of this document.



Technical Components

The functional requirements outlined in the architectural services section provide you with a general understanding of Vault's capabilities. Additional implementation-specific issues are critical to provide the best quality features. For example, functional scalability creates several technical issues: network bandwidth, catalog sizing, administration, etc. In this section, we list the specific components which deliver the architecture, and review how each component is designed to overcome various technical issues.

Components for Vault

- ◆ Vault Batch Processing
- ◆ Vault Duplication
- ◆ Vault Duplication Monitoring
- ◆ Vault Configuration
- ◆ Vault Reporting
- ◆ Vault Media Management
- ◆ Backup Image Duplication
- ◆ Existing NetBackup Services
- ◆ Media Manager Interface

Technical Design Issues

1. Initiate and control duplications from a centralized location
2. Initiate and control restores of duplicated image from a centralized location
3. Control various duplication parameters by use of NetBackup policies.
4. Reduce duplicating data over the network by duplicating locally whenever possible.
5. Support automated retry of duplication.
6. Support soft shutdown of duplication.
7. Direct appropriate duplications to either master or media servers.
8. Support one or more pairs of drives per server.

- 9. Write duplicates to a variety of storage devices and media.
- 10. Allow duplicates to span multi-volume media, yet support industry standard tar format for disaster recovery.
- 11. Create duplicates that can de-multiplex NetBackup images.
- 12. Interact with the Media Management service for media availability and media mount/unmount.
- 13. Create duplicate restores that work the same as normal Backup/Restore, for example, are allowed to the same or a different client, the same or a different location.
- 14. Allow duplications to notify external operations, for example, by integrating with the Event Management Service.
- 15. Support all normal Backup/restore functions supported, for example: data types, client types.

Vault Technical Components

This table lists the various steps as shown in the Example #1 diagram. Numbers in the table correspond to numbers in the diagram.

Vault Technical Components		
Service	Components	Technical Design
Vault Batch Functional Capability: Organize various steps into daily/weekly batch	vltrun, vltcore.acs, vltcore.tld	<ul style="list-style-type: none">Runs Vault utilities for each specific step.Simple script logicUses one script for ACLS processing, another for TLD processing.



<p>Vault Duplication Functional Capability:</p> <ul style="list-style-type: none"> Find all images to duplicate only for policies configured. Run one or more duplicates simultaneously on one or more servers 	<p>Vault - function preview and function duplicate</p>	<ul style="list-style-type: none"> Run <code>bpimagelist</code> for a given date range. <code>vltcore</code> will filter images based on policies, schedules, schedule types, media servers, and clients. Load balance duplications by splitting found backup images into files, one for each server/drive pair. Run multiple <code>bpduplicate</code> in parallel and monitors <code>bpduplicate</code> processes and log files. Run <code>bpduplicate</code> for each batch of images (based on media ID) found on server where duplicate was made if possible. (If backup server is one of the servers used for duplication.) Otherwise, run duplication on any available server. Can bypass duplication step to only vault original images, original media.
<p>Vault Duplication Monitoring Functional Capability: Monitor duplicates</p>	<p><code>vltcore</code> - function duplicate</p>	<ul style="list-style-type: none"> Monitor log file from <code>bpduplicate</code> for errors in background. Ensure each <code>bpduplicate</code> drive pair output is monitored separately and has unique identifier for support. Ensure <code>bpduplicate</code> has exited before allowing next <code>bpduplicate</code> to run on drive pair.
<p>Vault Configuration Functional Capability: Configures which backups to vault</p>	<p>Vault configuration file (<code>db/vault/ vault.xml</code>)</p>	<p>The <code>.xml</code> file read by <code>vltcore</code> contains:</p> <ul style="list-style-type: none"> server(s), number of drive pairs, destination storage unit, policies, schedules, schedule types, clients, media servers, date range, name of vault, NetBackup pools for vaulting and optional duplication, Netbackup robotic and non-robotic group, NetBackup catalog backup pool and retention period.

<p>Vault Report</p> <p>Functional Capability:</p> <p>Reports for picking used media from robot/library, for returning expired media from vault, for full vault inventory, for recovery.</p>	<p>vltcore - function report</p>	<ul style="list-style-type: none"> Reports use data files created by Vault Media Management commands Report “commands” are picking_library, dist_vault, dist_library, picking_vault, vault. (For more information, see “Reporting” on page 109.)
<p>Vault Media Management</p> <p>Functional Capability:</p> <p>Commands to find current Media Manager inventory</p>	<p>vltcore - command media</p>	<ul style="list-style-type: none"> Media function volumeinv - saves Media Manager data for vaulting pools
<p>Backup Image Duplication</p> <p>Functional Capability:</p> <p>Find images to duplicate</p> <p>Duplicate images</p>	<ul style="list-style-type: none"> bpduplicate 	<ul style="list-style-type: none"> Generates list of images to duplicate Sends command to bptm to copy image, location of log file. Runs in foreground. Waits for bptm to finish.
<p>NetBackup Backup Server</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> Master server listen for duplication requests Allow user-directed backups 	<p>VERITAS NetBackup Server utilities:</p> <ul style="list-style-type: none"> bprd bpbrm bptm bpdm 	<ul style="list-style-type: none"> Support network duplications. Duplicate True Image Recovery (TIR) separately for tracking file deletion. Maintain file / OS security Support multi-tape backup images
<p>NetBackup Master / Media Server</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> Master server sends duplicate request to appropriate server Provide central backup/duplicate image catalog for restore. 	<p>VERITAS NetBackup Master/Media Server:</p> <ul style="list-style-type: none"> bprd bpcd add_slave add_slave_on_clients 	<ul style="list-style-type: none"> Redirect duplication requests to appropriate Media Server based upon destination storage requested for specific image. Allows duplicate to run off-network for tape-to-tape copy.



<p>NetBackup Server side Multiplexing (MPX)</p> <p>Functional Capability:</p> <p>Support copy of tape multiplexing both on local backups and across networks.</p>	<p>VERITAS NetBackup server utilities:</p> <ul style="list-style-type: none"> • <code>bpbrm</code> • <code>bptm</code> 	<ul style="list-style-type: none"> • De-multiplexing of backup image automatic by <code>bptm</code>, and optionally preserve multiplexing information. • Only one duplicate per drive pair. • Duplication speed limited to tape read/write speeds. High use of multiplexing limits copy speed to tape read and tape mount for all tapes needed for de-multiplexing, assuming tape read/write speeds similar.
<p>NetBackup Backup Server Database Services</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> • Provides Image information including number of copies. • Stores media used information for both primary copy and duplicate copies. 	<p>VERITAS NetBackup database daemon/service:</p> <ul style="list-style-type: none"> • <code>bpdbm</code> 	<ul style="list-style-type: none"> • Use <code>bpdbm</code> for all data access requests • <code>bpdbm</code> daemon/service process always running on master server - not on media servers. • <code>bpdbm</code> creates unique backup identifier based on client name, time of day. • <code>bpdbm</code> provides image catalog, number of image copies for determining which images need duplication. • <code>bpdbm</code> provides backup media catalog for determining which backup media contains backup image, and which duplication media is used for copies. • <code>bpdbm</code> provides restore with duplicate copy image/media when "primary copy" is not 1.
<p>NetBackup Schedules - Centralized server scheduling</p> <p>Functional Capability:</p> <p>Can be used by <code>bpduplicate</code> to limit duplication of a policy to only images within specific schedule</p>	<p>VERITAS NetBackup Scheduler:</p> <ul style="list-style-type: none"> • <code>bpsched</code> 	<p><code>vltdcore</code> preview can be limited by schedule name or type.</p>

<p>NetBackup Backup Server - Restore functions</p> <p>Functional Capability: Works as normal for duplicated images.</p>	<p>VERITAS NetBackup server daemons/processes:</p> <ul style="list-style-type: none"> ◆ bprd ◆ bpcd ◆ bpbrm ◆ bptm ◆ bpdm 	<ul style="list-style-type: none"> ◆ Duplicated images can be de-multiplexed. ◆ Changing the primary copy automatically forces any restore of that image to use the specified primary copy. ◆ There is currently no GUI interface to select the duplicated copy. The change requires manual intervention.
<p>NetBackup policies - Centralized client management</p> <p>Functional Capability: Works as normal. Policy name used to limit duplication to specific clients.</p>	<p>VERITAS NetBackup policy utilities</p>	<ul style="list-style-type: none"> ◆ Name of clients stored in the configuration file.
<p>NetBackup Storage Unit - Centralized resource control</p> <p>Functional Capability: Works as normal. Used to determine which server will run duplicate</p>	<p>VERITAS NetBackup storage unit utilities</p>	<ul style="list-style-type: none"> ◆ Destination storage unit in configuration file will match up with source media server.



<p>NetBackup Image Catalog Services - centralized image information</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> Centrally stores backup and duplicate catalog information. Keeps track of both backup and duplicate image information: media, fragments, size, location on tape required for recovery. Lists directories/files/fragments in image catalog. Tracks backup and duplicate server name for each image to support master/media server duplication 	<p>VERITAS NetBackup image catalog utilities:</p> <ul style="list-style-type: none"> bptm bpimage bpimagelist bplist bpflist bpimmedia bpfrag 	<p>The master server stores all image data. Duplicates notify the master server of copy status - fragments, media used, but no image information necessary.</p>
---	--	---

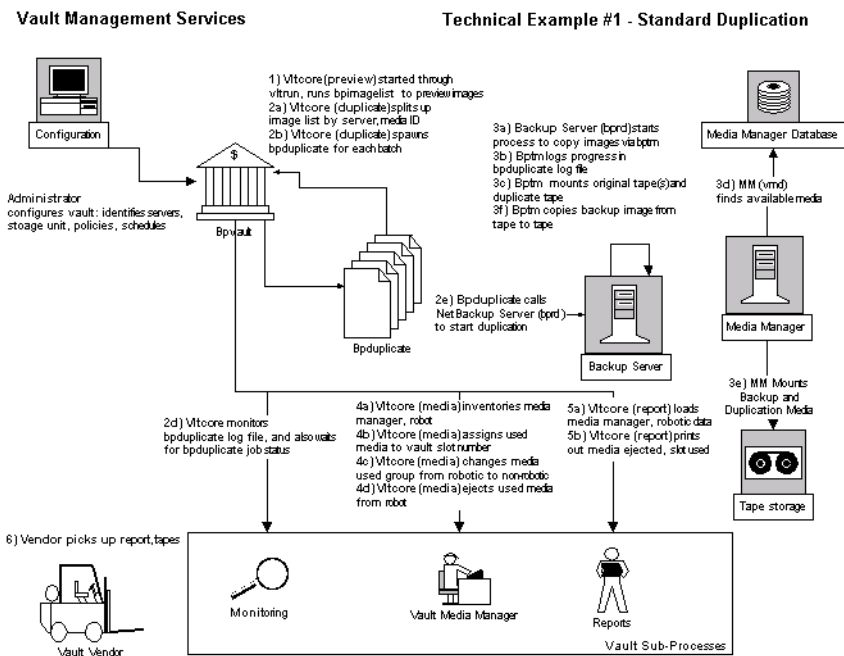
<p>NetBackup Media Manager - centralized media control</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> ♦ Used to determine which media used for primary backup copy. ♦ Can manually expire primary copies media to force use of duplicate copy. ♦ Backup retention periods for primary backup image can differ for duplicate image. ♦ Expire media after retention period will expire duplicates copies. ♦ Allow freeze media sent off-site to keep from expiring while in need for recovery. 	<p>VERITAS NetBackup media database utilities:</p> <ul style="list-style-type: none"> ♦ bptm ♦ bpmedia ♦ bpmedialist ♦ bpexpdate ♦ bplabel 	<ul style="list-style-type: none"> ♦ Media used for image duplicates can have a different retention period. ♦ Duplicate media used must be assigned in Media Manager to duplicate pool named in configuration file. ♦ Expired duplicate media released in Media Manager triggers recall from off-site vault. ♦ Duplication sends image fragment, media used information to master server.
<p>Backup monitoring</p> <p>Functional Capability:</p> <ul style="list-style-type: none"> ♦ Duplication monitored by NetBackup monitoring tools. ♦ Vault log files can be used to track status. 	<p>Activity Monitor - VERITAS NetBackup daemon/service logging.</p>	<ul style="list-style-type: none"> ♦ VERBOSE option increases information to logging subdirectories. ♦ Notify scripts supported for start and end session, and for pre-and post-eject processing.
<p>Backup reports</p> <p>Functional Capability:</p> <p>Provide administrative reports needed for capacity, utilization planning, auditing purposes, technical support.</p>	<p>See Backup Console.</p> <ul style="list-style-type: none"> ♦ bpterror ♦ bpimagelist ♦ cleanstats ♦ available_media ♦ support 	<ul style="list-style-type: none"> ♦ Can load backup information into RDBMS for better reporting. ♦ bpimagelist provides basic information showing all copies of an image.



<p>Media Manager interface</p> <p>Functional Capability:</p> <ul style="list-style-type: none">♦ Use Media Manager Tools to request media from specific pools; mount tapes; control robotics♦ Use specific Media Manager pool for duplication.		<ul style="list-style-type: none">♦ Must ensure NetBackup Media database and Media Manager database stay synchronized.♦ Update Media Manager fields to store vault name, duplication session, slot id, date requested, date sent off site.
---	--	---



Technical Example: Standard Duplication Diagram



Technical Example: Standard Duplication Table

Technical Example #1 Standard Duplication/Vault		
Service	Components	Outgoing Program/Data Flow
<p>Backup Server-File System and Raw Partition</p> <p>Incoming Program/Data Flow:</p> <ul style="list-style-type: none">• bprd started by startup script on master server via command line. bprd must always be running to allow any backup/restore command.• bprd can be started by a GUI Initiate Request Daemon option• bprd called by bpbackup, bpduplicate, bprestore, bparchive to start jobs. Uses known socket in /etc/services• Debug Log: /usr/open/ netbackup/logs/ bprd on server	<p>bprd - Job request daemon/service.</p>	<ol style="list-style-type: none">1. bprd starts bpsched on master server via command line interface to handle jobs. <ul style="list-style-type: none">• For manual jobs, bprd builds file list from client (for example, bpbackup), then bprd starts bpsched as normal.• bprd regularly cleans up debug logs.



<p>Schedules - Centralized backup scheduling</p> <p>Incoming Program/Data Flow</p> <p>Scheduler started by <code>bprd</code> on master server.</p> <p>Scheduler exits if no jobs needed to run, or monitored client job is finished.</p> <p>Debug Log: <code>/usr/opensv/ netbackup/logs/ bpsched</code> on Master server</p>	<p>Master Server Scheduler: <code>bpsched</code></p>	<ul style="list-style-type: none"> • <code>bpsched</code> calls <code>bpdbm</code> to obtain policy information, storage unit server to use, file list, etc. • <code>bpsched</code> starts <code>bptm</code> to do basic media check prior to starting client job. • <code>bpsched</code> starts <code>bpcd</code> via <code>inet.d</code> on appropriate storage unit server (e.g. master or media server). • <code>bpsched</code> sends <code>bpbrm</code> command for <code>bpcd</code> to run, with required options, file list, log path to <code>bpcd</code> on server via socket • <code>bpsched</code> monitors <code>bpbrm</code> output via <code>stderr</code> • <code>bpsched</code> signals other <code>bpsched</code> processes to ensure only one scheduler is acting as main scheduler, to eliminate accidental duplication of jobs.
<p>See Example #1 Diagrams</p> <p>Incoming Program/Data Flow</p> <p><code>bpcd</code> started by <code>bpsched</code> via <code>inetd</code> / known socket. Acts as job proxy for <code>bpsched</code></p> <p>Debug Log: <code>/usr/opensv/ netbackup/logs/bpcd</code> on server for server related commands</p>	<p><code>bpcd</code> Backup job proxy</p>	<p><code>bpcd</code> starts backup/restore manager <code>bpbrm</code> on master or media server</p>



<p>Backup Server (cont'd)</p> <p>Incoming Program/Data Flow</p> <ul style="list-style-type: none">◆ bpbrm started by bpsched on storage unit server (via bpcd). One bpbrm process started for each backup or restore operation.◆ bpbrm stderr output and exit status monitored by bpsched◆ bpbrm exits to bpsched with status◆ bpbrm receives signal() from bptm when media ready.◆ bpbrm sends email on job completion (via bpcd) <p>Debug Log: /usr/opensv/ netbackup/logs/ bpbrm on server</p>	<p>bpbrm Backup/Restore Manager</p>	<ul style="list-style-type: none">◆ bpbrm starts bptm to start backups and duplication. Waits for child exit status.◆ bpbrm starts bpbkar on client. Waits for child exit status.
---	---	--



<p>Backup Server (cont'd) Incoming Program/ Data Flow</p> <ul style="list-style-type: none"> • bptm started by bpsched to check validity of storage unit. • bptm started by bpbrm on master and Media Servers via command line. One bptm started for backup/restore when using tape or optical media. • Child bptm started by Parent bptm. • bptm messages bpbrm once media is mounted via signal. • bptm stderr output monitored by bpbrm • Child bptm exits to parent bptm on client completion. • Parent bptm exits to bpbrm on server completion. <p>Debug Log (on server): UNIX: <code>/usr/openv/netbackup/logs/bpsched</code> Windows: <code>install_path/netbackup/logs/bpsched</code></p>	<p>bptm: Server tape manager</p>	<ul style="list-style-type: none"> • Parent bptm starts child bptm. • Parent bptm calls vmd on master server via known socket to find appropriate media. • Parent bptm also calls bpdbrm to compare Media Manager database with own, NetBackup Media database. • Parent bptm calls ltid on storage unit server via known socket to mount media. • Child bptm receives data from bpbkar on client. • Child bptm writes data to buffer. • Parent bptm writes buffers to tape when they are full. <p>On backup completion, parent bptm runs:</p> <p>UNIX: <code>/usr/openv/netbackup/bin/backup_notify</code></p> <p>Windows: <code>install_path/netbackup/bin/backup_notify</code></p>
--	----------------------------------	---



<p>Backup Server Policy, Image and Media Database Services</p> <p>Incoming Program/Data Flow:</p> <ul style="list-style-type: none"> • bpdbrm started by initbprd during OS boot on master server only. • In this example, bpdbrm called by bpsched, bpbbrm, bptm. • bpdbrm daemon/service process always running on master server - not on media servers. • bpdbrm database can be accessed locally or across network using known socket for bpdbrm. • bpdbrm creates unique backup identifier based on client name, time of day. 	<p>Master Server database daemon/service: bpdbrm</p>	<ul style="list-style-type: none"> • bpdbrm provides policy data to bpsched to build worklist. • bpdbrm provides configuration information to bpbbrm. • bpdbrm provides bptm with information about Backup Media database for comparison with Media Manager database. • bpdbrm stores Media used information into Backup Media database for bptm • bpdbrm stores Image file list catalog for bpbbrm.
<p>Backup Media Manager</p> <p>Incoming Program/Data Flow:</p> <p>Backup Media management handled by calls from bptm to bpdbrm shown above.</p>	<ul style="list-style-type: none"> • bptm • bpdbrm 	<p>Parent bptm calls bpdbrm to save media used data to media database. Updates keyword search for user-directed backups.</p>
<p>Backup Image catalog</p> <p>Incoming Program/Data Flow:</p> <p>Image catalog handled by calls from the client bpbkar to server bpbbrm and then from bpbbrm to bpdbrm.</p>	<ul style="list-style-type: none"> • bpbkar • bpbbrm • bpdbrm 	<ul style="list-style-type: none"> • bpbkar writes image catalog data to bpbbrm on server. • bpbbrm calls bpdbrm via socket to store image file list catalog.

<p>Backup Client</p> <p>Incoming Program/Data Flow:</p> <ul style="list-style-type: none">• bpcd started by bpbrm (via inetd) to start backup jobs. Acts as proxy for bpbrm.• bpcd can perform other chores, e.g. sends mail for bpbrm.• bpcd returns exit status to bpbrm via socket. <p>Debug Log:</p> <p>UNIX:</p> <p><i>/usr/opensv/ netbackup/logs/ bpcd</i></p> <p>Windows:</p> <p><i>install_path/ netbackup/logs/bpcd</i></p>	<p>bpcd - client job proxy</p>	<p>bpcd starts bpbkar on client for backup.</p>
---	--------------------------------	--



<p>Backup Client, continued Incoming Program/Data Flow:</p> <p>bpbkar started by bpbm (via bpcd).</p> <ul style="list-style-type: none"> • bpbkar receives backup policy options, file list, etc. from bpbm. This information was originated by bpsched at step #1. • bpbkar exit status is sent to bpsched via socket. <p>Debug Log:</p> <p>UNIX:</p> <p>/usr/openv/ netbackup/logs/ bpbkar</p> <p>Windows:</p> <p>install_path/ netbackup/logs/ bpbkar</p>	<p>bpbkar - client data transfer</p>	<ul style="list-style-type: none"> • bpbkar reads file/directory list from bpbm. • bpbkar determines NFS mounts and adds or ignores depending on policy. • bpbkar compresses client data if necessary. • bpbkar writes client data to bptm/bpdm on server.
<p>Media Manager Incoming Program/Data Flow:</p> <ul style="list-style-type: none"> • Media Manager vmd called by server tape manager bptm for scratch tape. • Media manager ltid called by server tape manager bptm to mount tape. <p>Note See Media Manager functional design for detailed information.</p>	<ul style="list-style-type: none"> • vmd • ltid 	<ul style="list-style-type: none"> • Media Manager provides new media ID for backup. Uses pool for determining which media is valid. • Media Manager ltid mounts tape.

Operational Procedures

This table summarizes Operational Procedures for Vault. The NetBackup Vault 4.5 *Operator's Guide* provides more detailed information on day to day procedures. The NetBackup Vault 4.5 *System Administrator's Guide* provides more detailed information on installation, configuration and troubleshooting.

Vault Management / Operational Procedures		
Service	Operational Procedure	Staff Responsibilities
Vault Configuration	<ol style="list-style-type: none"> 1. Review backup procedures; determine duplication capacity needed. 2. Assign appropriate server to run duplications; determine appropriate window for running duplication. 3. Configure Vault parameters via Administration Console. 4. Review duplication windows for performance, throughput. 	<ul style="list-style-type: none"> ◆ Determine need for basic levels of duplication service on a per policy basis. ◆ Ensure sufficient hardware, software, network capacity is available for duplication of backup images.
Vault Duplication	Set up Vault policy to schedule vault sessions.	Start duplication job on time, daily and/or weekend.
Vault Monitoring	<ol style="list-style-type: none"> 1. Use Activity Monitor to determine progress. 2. Set up links between log file and monitoring system for email and/or paging notification. 	<ul style="list-style-type: none"> ◆ Ensure duplication jobs complete successfully. ◆ Ensure errors reported to Event Management.
Vault Report	Compare report output with ejected tapes, returned tapes	Ensure accuracy of vault process.
Vault Media Management	Check duplication volume pools and catalog backup pools for sufficient media.	Ensure sufficient media available for duplication.



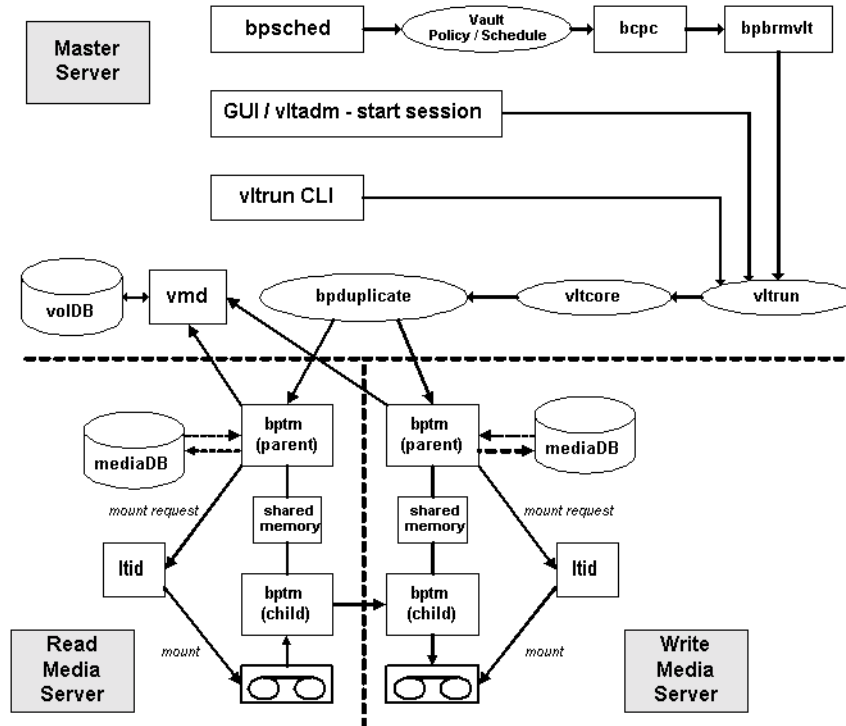
Backup Media Manager	<ol style="list-style-type: none"> 1. Use Media Manager to manage media. 2. Manually expire/freeze tapes when needed for retrieval from vault. 	
Backup Image Catalog	<ol style="list-style-type: none"> 1. Set up schedule for backup of image catalog. 2. Ensure specific tapes available to store catalog. 	Ensure backup catalog is safely copied.
Event Management Interface	Set up appropriate Event response procedures.	Same as for normal Backup.
Incident Management Interface	Set up procedure for passing storage events to Help Desk.	Same as for normal Backup.
Backup Reports	Run regular reports to ensure proper images are duplicated.	Review production duplication cycle for thoroughness.
Duplication Capacity Review	Determine capacity planning cycle, including reaction time, costing factors, and new requirements.	<ul style="list-style-type: none"> ◆ Assist production support over time on determining system, robotic, network utilization rates and effective valuation, for example, disk capacity. ◆ Assist in new requirements and performance related additions to the system infrastructure.
Recovery Review	Run regular tests to ensure recovery of essential data from off-site storage.	<ul style="list-style-type: none"> ◆ Ensure knowledge of appropriate procedures for restoring duplicated images. ◆ Ensure sufficient knowledge to restore database catalog, backup software, etc. in case of disaster on Netbackup server(s).

Media Manager	<ol style="list-style-type: none">1. Determine media requirements and setup initial media pool for duplication.2. Monitor media pool usage over time.3. Configure master/media server media management.	Ensure sufficient media is available for duplicates to run.
---------------	---	---



NetBackup Vault Image Duplication Process

The diagram below shows the flow of operations for the duplication process during a vault session. For simplicity, this diagram shows duplication from tape-to-tape, and shows operations for creating a single copy.



A vault session is initiated by:

- ◆ The NetBackup scheduler (**bpsched**) finding a Vault schedule with an open window.
- ◆ The operator issuing a Start Session for a particular Vault profile from the NetBackup Administration Console – Vault Management GUI or from the **vltadm** MUI.
- ◆ The operator issuing a **vltun** command from the command line.

Regardless of the method that initiated the vault session, it will result in **vltun** being started. These are the stages of the process:

Process	Actions
vltrun	vltrun calls vltcore. vltcore starts bpduplicate. vltcore does quite a bit of pre-processing to determine what images need to be vaulted/duplicated, etc.
bpduplicate	bpduplicate starts bptm on the READ SERVER and on the WRITE SERVER.
bptm	bptm on the READ SERVER and on the WRITE SERVER respectively locate media to be used in the duplication process (lookups to mediaDB and volDB) and issue mount requests (tpreq) for those media.
bptm (child process)	bptm (child process) read the data from the READ SERVER media, send it to the WRITE SERVER where the bptm (child process) writes the data to the destination media





Glossary

access control list (ACL)

Security information associated with files on some file systems.

ACS

Automated Cartridge System. ACS can refer to any of the following:

- ◆ A type of Media Manager robotic control. This robot type is supported only by NetBackup DataCenter servers.
- ◆ The StorageTek (STK) system for robotic control.
- ◆ The highest-level component under STK's ACS library software, which refers to a specific standalone robotic library or to multiple libraries connected with a media passthru mechanism.

active job

A job for which NetBackup is currently processing backup or restore data.

activity logs

See "debug logs."

activity monitor

A NetBackup administration utility that displays information about NetBackup jobs and provides limited control over them.

administration client

See "remote administration console."

administrator

A user that is granted special privileges to install, configure, and manage the operation of a system, network, or application



AIT

Sony Advanced Intelligent Tape, a type of tape drive or media type.

alternate-client restore

See “redirected restore (different client).”

alternate-target restore

See “redirected restore (different target).”

alternate path restore

See “redirected restore (different path).”

alternate read server

A server used to read a backup image which was originally written by a different media server. The media server specified as Alternate Read Server must have access to the media containing the backup image or images it is configured to read.

archive

A special kind of backup where NetBackup backs up the selected files, and if the backup is successful, deletes the files from the local disk. In this manual, references to backups also apply to the backup portion of archive operations except where otherwise noted.

archive bit

A file-status bit that the Microsoft based operating system sets when it writes a file, thereby indicating that the file has changed.

attributes for a policy

Configuration parameters that control the behavior of NetBackup during operations involving this policy.

autochanger

See “robotic library.”

autoloader

See “robotic library.”

automatic backup

A scheduled backup by the master server.

back up

The act of copying and saving files and folders to storage media.

backup

Refers to the process of copying and saving files and directories to storage media. For example, *the backup is complete*. This term can also refer to the collection of data that NetBackup saves for a client during a backup or archive. For example, *duplicate the backup*.

Backup is two words when used as a verb. For example, *back up the file*.

backup, archive, and restore interface

The name of the NetBackup Microsoft Windows and Java based user interfaces for clients. On servers, these interfaces can be started through the NetBackup Administration Console.

backup window

The period of time during which backups can begin.

block size

The number of bytes in each block of data written on the media during a backup.

bp

A backup, archive, and restore utility for users on NetBackup UNIX clients. It has a character-based, menu interface that can be run from terminals that do not have X Windows capabilities.

bpadm

An administrator utility that runs on NetBackup UNIX servers. It has a character-based, menu interface that can be run from terminals that do not have X Windows capabilities.

bp.conf file

A NetBackup configuration file on UNIX servers and also on UNIX, Macintosh, and OS/2 clients.

bp.ini file

NetBackup initialization file for Novell NetWare target clients.

bpcd

NetBackup Client service on Windows and the NetBackup Client daemon on UNIX.



bprd

NetBackup Request Manager service on Windows and NetBackup Request daemon on UNIX.

cancel a job

Terminating a job and removing it from the job queue.

carousel

See “robotic library.”

catalogs

Internal NetBackup and Media Manager databases. These catalogs contain information about configuration, media, devices, status, errors, and the files and directories in the stored backup images.

CDF

Context-dependent file, which is a type of directory structure on a Hewlett-Packard system.

changer

See “robotic library.”

class

See “policy.”

client

The system with the files to back up, archive, or restore.

client-user interface

See “user interface.”

cluster

See master and media server cluster.

command lines

Commands that users can execute either from the system prompt or in scripts.

compression

The process of compacting data to enable more efficient transmission and storage.



configuration

The parameters that govern the behavior of an application. This term can also refer to the manner in which a network or system is laid out or connected (for example, a network configuration).

consolidated eject

A process of ejecting media for more than one Vault session at a time. A Consolidated Eject can be performed for one or more logical vaults at one time.

consolidated report

A process of generating reports for more than one Vault session at a time. A Consolidated Report can be performed for one or more logical vaults at one time. Consolidated reports are organized by report title, not by vault.

cpio

A UNIX command that can be used for copying files to or from a cpio archive on disk or tape.

ctime

The time that a UNIX inode was changed.

cumulative-incremental backup

A backup that is scheduled by the administrator on the master server and backs up files that have changed since the last successful full backup. All files are backed up if no prior backup has been done. Also see “differential-incremental backup.”

daemon

A program on a UNIX system that runs in the background and performs some task (for example, starting other programs when they are needed). Daemons are generally referred to as services or processes on Windows server systems.

database-agent clients

Clients with additional NetBackup software that is designed to back up relational databases.

database-extension clients

See “database-agent clients.”



debug logs

Logs that can be optionally enabled for specific NetBackup and Media Manager programs and processes and then used to investigate problems.

destination storage unit

A storage unit to which Vault sends the data from a duplication operation. If the duplicated backup images are to be vaulted, then the destination storage unit must correspond to the robotic volume group.

device delays

Delays caused by the device that are beyond the control of the storage application. An example is the time required to position tape under the read and write heads.

device host

A host (that has Media Manager installed) where a drive or robotic control is attached or is defined.

device monitor

A Media Manager administration utility that provides monitoring and manual control of Media Manager storage devices. For example, an administrator or computer room operator can use this utility to manually reset devices or set them to the UP or DOWN state.

DHCP

Dynamic host configuration protocol. This TCP/IP protocol automatically assigns temporary IP addresses to hosts when they connect to the network.

differential-incremental backup

Scheduled by the administrator on the master server and backs up files that have changed since the last successful incremental or full backup. All files are backed up if no prior backup has been done. Also see “cumulative-incremental backup.”

directory depth

The number of levels below the current directory level that the NetBackup interfaces show in their directory and file list displays.

directory tree

The hierarchical structure in which files are organized on a disk. Each directory lists the files and directories that are directly below it in the tree. On UNIX, the topmost directory is called the root directory.

disaster recovery

Recovering data from backups after a disk crash or other catastrophe.

disk

Magnetic or optical disk storage media.

disk-image backup

A bit-by-bit rather than a file system backup of a disk drive on a Windows platform.

DLT

Digital-linear tape or tape drive type.

Domain Name Service (DNS)

A program that handles name translation for network communications.

drive cleaning

The use of a special cleaning tape to clean the heads on a drive.

duplicate image

A copy of a backup image.

eject

Move media out of a robotic library.

encryption

Provides additional security by encrypting backup data on the client. This capability is available only with the NetBackup Encryption option.

entry and exit ports

See “media access port.”

exclude list

A list that designates files or directories to exclude from automatic backups.

expiration (image)

The date and time when NetBackup stops tracking a backup image.



expiration (volume)

The date and time when the physical media (tape) is considered to be no longer usable.

external media ID

This is an identifier written on a media cartridge or canister to help the operator identify the volume before inserting it into a drive or robot. For labeled media, the external media ID should be the same as the media ID recorded on the media.

EVSN

See “external media ID.”

FlashBackup

A special type of raw-partition backup that requires the NetBackup FlashBackup separately-priced option (this option is available only for NetBackup DataCenter).

flush level

Controls how often Netbackup clears its log files on a Novell NetWare or Microsoft Windows client platform.

fragment

A part of a backup or archive image. NetBackup can be configured to divide images into fragments when they exceed a certain size or span tapes.

frequency (backup)

How often NetBackup performs scheduled backups. For example, if the frequency is seven days then backups occur once a week.

FROZEN media state

If a volume is FROZEN, NetBackup keeps it indefinitely and can restore from it but not use it for further backups or archives.

full backup

A backup that copies, to a storage unit, all files and directories that are beneath a specified directory.

FULL media state

If this appears in a report or listing, it indicates the volume is FULL and cannot hold more data or be used for further backups.

global attributes

NetBackup configuration attributes that affect all policies.

GDM Dashboard

The name for the Global Data Manager interface. The Dashboard enables monitoring job and drive activity on multiple master servers, as well as providing alerts to problem conditions.

GDM Managed Server

A NetBackup master server that appears as a managed master server in the left pane of the GDM Dashboard.

GDM Server

A NetBackup master server that has the Global Data Manager license activated. When logging into this host, the user can monitor the activity on multiple master servers using the GDM Dashboard interface. If the host has installed the Advanced Reporter option, the reports show information on multiple master servers.

Global Data Manager (GDM)

A separately-priced option (for UNIX servers) that provides an interface with a tree view where the administrator can view and administer multiple master servers. The server where the option is installed is called a GDM Server.

Global Device Database

A single host that serves as the repository for global device configuration information. When you install NetBackup, by default the master server is configured as the global device database host.

GNU tar

A public domain version of the UNIX tar program.

goodies directory

A directory containing programs, scripts, and other files that are not formally supported.

GUI

Graphical user interface.



hard link

On UNIX, a hard link is a pointer to the inode for the data. On a Windows server, a hard link is a directory entry for a file. Every file can be considered to have at least one hard link. On NTFS volumes each file can have multiple hard links, and a single file can appear in many directories (or even in the same directory with different names).

heap level

A parameter for memory-heap debugging on a Novell NetWare or Windows NetBackup client.

hierarchical storage management

The process of automatically migrating selected files from a managed file system to specified migration levels on secondary storage, while maintaining transparent access to those files.

host

A computer that executes application programs.

host name

Name by which a host computer is identified by programs and other computers in the network.

HSM

See storage migrator.

image

The collection of data that NetBackup saves for an individual client during each backup or archive. The image contains all the files, directories, and catalog information associated with the backup or archive.

import

The process of recreating NetBackup records of images so the images can be restored.

include list

A list that designates files or directories to add back in from the exclude list.

incremental backup

See “cumulative-incremental backup” and “differential-incremental backup.”

inject

Move media into a robotic library.

inport

See “media access port.”

inode

A UNIX data structure that defines the existence of a single file.

install_path

Directory where NetBackup and Media Manager software is installed. The default on Windows servers is C:\Program Files\VERITAS and on UNIX it is /usr/openv.

jbpSA

The Java-based NetBackup interface for performing user backups, archives, and restores.

jnbSA

The Java-based NetBackup interface for administrators.

job

A parcel of work submitted to a computer. NetBackup jobs are backups, archives, or restores.

kernel

The nucleus of an operating system.

keyword phrase

A textual description of a backup.

kill a job

See “cancel a job.”

label

Identifier of a tape or optical disk volume. A recorded label includes a media ID.
A barcode label allows a barcode scanner to be used for media tracking.

library

See “robotic library.”



link

See “hard link” or “symbolic link.”

LMF - Library Management Facility

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

This robot type is supported only by NetBackup DataCenter servers.

load

(noun) Amount of work that is being performed by a system or the level of traffic on a network. For example, network load affects performance.

(verb) Copy data to internal memory. For example, load the installation program.

(verb) Used to indicate tape drive initialization done when new media is being added.

logs

Files where a computer or application records information about its activities.

mailslot

See “media access port.”

man pages

Online documentation provided with UNIX computer systems and applications.

Master and media server cluster

A NetBackup master server and the remote media servers that it is using for additional storage. It is possible to configure clusters only with NetBackup DataCenter servers. NetBackup BusinessServer supports only a single server, the master.

Master of Masters

A NetBackup host where Global Data Manager software is installed. When logging into this host, the interface has a tree view where the administrator can view and administer multiple master servers.

master server

The NetBackup server that provides administration and control for backups and restores for all clients and servers in a master and media server cluster. NetBackup BusinessServer supports only a single server and it is the master.

media

Physical magnetic tapes, optical disks, or magnetic disks where data are stored.

media access port

A slot or other opening in a robot where you can insert or remove a tape without having to access the interior of the robot. After inserting a tape, you move it to a slot by using an inject command. Prior to removing a tape, you move it to the port by using an eject command. The inject and eject commands are supported through the add and move screens in the Media Manager administration interface.

media host

NetBackup server to which the job (client) is sending the data.

media ID

An identifier that is written on a volume as part of the recorded label.

Media Manager

Software that is part of NetBackup and manages the storage devices and removable media.

Media Manager Host

Host where Media Manager is installed (may have devices attached)

media server

A NetBackup server that provides storage within a master and media server cluster. The master can also be a media server. A media server that is not the master is called a remote media server. NetBackup BusinessServer does not support remote media servers.

menu interface

A character-based interface for use on terminals that do not have graphical capabilities.

mount

Make a volume available for reading or writing.

mount point

The point where a file system on a disk logically connects to a system's directory structure so the file system is available to users and applications.



MPX

See “multiplexing.”

mtime

The point in time when a UNIX or NTFS file is modified.

multiplexing

The process of sending concurrent-multiple backups from one or more clients to a single storage device and interleaving those images onto the media.

multiplexed group

A set of backups that were multiplexed together in a single multiplexing session.

NDMP

Network data management protocol. NetBackup requires the NetBackup for NDMP separately-priced option to support NDMP.

NetBackup Client service

NetBackup Windows service that runs on clients and servers and listens for connections from NetBackup servers and clients in the network. When a connection is made, this service starts the necessary programs.

NetBackup configuration options

On UNIX servers and on UNIX and Macintosh, clients, these settings are made in the `bp.conf` file. On NetWare target and OS/2 clients, they are in the `bp.ini` file. On Windows servers and Windows clients, these settings are called properties and are made through the Backup, Archive, and Restore interface or the Host Properties dialog in the NetBackup Administration Console.

NetBackup databases

See catalogs.

NetBackup Database Manager service

NetBackup Windows service that runs on the master server and manages the NetBackup internal databases (called catalogs). This service must be running on the master server during all NetBackup administrative operations.

NetBackup Device Manager service

The NetBackup Windows service that runs on a NetBackup server and starts the robotic control processes and controls the reservation and assignment of volumes. This service runs only if the server has devices under Media Manager control. The process is `ltid`.

NetBackup properties

Same as NetBackup configuration options but are called NetBackup properties on Microsoft Windows platforms.

NetBackup Request Manager service

The NetBackup Windows service that runs on the master server and starts the scheduler and receives requests from clients.

NetBackup Volume Manager service

A NetBackup Windows service that runs on a NetBackup server, allows remote administration of Media Manager, and manages volume information. The process is `vmd`.

NIS

Network information service.

NLM

NetWare loadable module.

NFS

Network file system.

nonrobotic

See “standalone.”

ODL

Optical disk library. This robot type is supported only by NetBackup DataCenter servers.

offsite volume group

A volume group in which media will appear after having been ejected from the robot for vaulting. When Vault ejects media it is moved from the robotic volume group to the off-site volume group.



offsite volume pool

A volume pool that contains media that is to be ejected and vaulted. Backup images written to an off-site volume pool by an original NetBackup backup policy or by Vault's duplication feature will be ejected and vaulted. More than one off-site volume pool can be specified for the Eject step of a Vault profile.

original backup

A backup image created by a backup job. A single backup image or all backup images created by an Inline Tape Copy (multiple copy) configuration are considered original backups. A backup image created by a duplication job is not an original backup.

outport

See “media access port.”

partitions

The logical partitions into which a magnetic disk is divided.

patch

A program that corrects a problem or adds a feature to an existing release of software.

path length

Number of characters in a pathname.

pathname

The list of directories in the path to a destination directory or file.

PC clients

NetBackup clients that have Microsoft Windows, Macintosh, or IBM OS/2 operating systems.

peername

The name by which a computer identifies itself when establishing connections to other systems.

policy

Defines the backup characteristics for a group of one or more clients that have similar backup requirements.

port

A location used for transferring data in or out of a computer.

Also see “media access port.”

primary copy

The copy of an image that NetBackup uses to satisfy restores. When NetBackup duplicates an image, the original is designated as the primary copy.

privileges

The tasks or functions that a user, system, or application is authorized to perform.

profile

A vault profile is a way to save configuration settings. Specific parameters for duplication, catalog backup, eject, and report or any combination of these steps, are configured within a profile.

progress report

Log where NetBackup records events that occur during user operations.

proxy restore

A proxy restore allows the user to restore files that he has write access to, on a machine other than his desktop. The files must be in a backup of the machine to which they are being restored.

QIC

Quarter-inch-cartridge tape.

queued job

A job that has been added to the list of jobs to be performed.

raw-partition backup

Bit-by-bit backup of a partition of a disk drive on UNIX. On Windows, this is called a disk-image backup.

rbak

The program that Apollo clients use to read data from tape during a restore.



recorded media ID

This is an identifier written as part of the label on a volume and used by Media Manager to ensure that the correct volume is mounted. The recorded media ID should match the external media ID.

redirected restore (different client)

Restoring files to your client when they were originally backed up from a different client. The administrator using the interface on the master server can direct a restore to any client (this variation is called a server directed restore).

redirected restore (different target)

On a Novell NetWare server platform running the NetBackup target version of client software, this operation restores files to a different target than the one from which they were backed up.

redirected restore (different path)

Restores files to a different directory than the one from which they were backed up.

registry

A Microsoft Windows database that has configuration information about hardware and user accounts.

remote administration console

A Windows NetBackup client that has the administration interface software installed and can be used to administer NetBackup servers.

remote media server

A media server that is not the master. Note that only NetBackup DataCenter supports remote media servers. NetBackup BusinessServer supports only a single server, the master.

residence

In Media Manager, information about the location of each volume is stored in a volume database. This residence entry contains information, such as robot number, robot host, robot type, and media type.

resource

A Novell NetWare term that refers to a data set on the target. For example, in DOS, resources are drives, directories, and files. Also see “target service.”

restore

(verb) The act of restoring selected files and directories from a previous backup or archive and returning them to their original directory locations (or to a different directory).

(noun) The process of restoring selected files and directories from a previous backup and returning them to their original directory locations (or to a different directory).

retention level

An index number that corresponds to a user-defined retention period. There are 10 levels from which to choose (0 through 9) and the retention period associated with each is configurable. Also see “retention period.”

retention period

The length of time that NetBackup keeps backup and archive images. The retention period is specified on the schedule.

robotic arm

The component of a robotic library that physically selects the media (tape or optical disk).

robotic library

Refers to a robot and its accompanying software. A robotic library includes a collection of tapes or optical platters used for data storage and retrieval. For example, a Tape Library DLT (TLD) refers to a robot that has TLD robotic control.

robotic volume group

A volume group from which media will be ejected and vaulted. When Vault duplicates backups, they are duplicated to media in the robotic volume group.

root

The highest level directory in a hierarchical directory structure. In MS-DOS, the root directory on a drive is designated by a backslash (for example, the root on drive C is C:\). On UNIX, the root directory is designated by a slash (/).

Also, a UNIX user name having administration capability.

RS-232

An industry-standard interface for serial communications and sometimes used for communicating with storage peripherals.



RSM Interface

Application in Windows 2000 used to manage Removable Storage Manager (RSM) devices.

RSM - Removable Storage Manager

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

Also, a component of the Windows 2000 operating system that manages storage devices.

RVSN

See “recorded media ID.”

schedules

Controls when backups can occur in addition to other aspects of the backup, such as: the type of backup (full, incremental) and how long NetBackup retains the image.

SCSI

Small computer system interface. This is a type of parallel interface that is frequently used for communicating with storage peripherals.

server-directed restore

Using the user interface on the master server to restore files to any client. Only the administrator can perform this operation.

server independent restore

Restoring files by using a NetBackup server other than the one that was used to write the backup. This feature is available only with NetBackup DataCenter.

server list

The list of servers that a NetBackup client or server refers to when establishing or verifying connections to NetBackup servers. On a Windows server and Microsoft Windows clients, you update the list through a dialog box in the interface. On a UNIX server and UNIX and Macintosh clients, the list is in the `bp.conf` file. On NetWare target and OS/2 clients, the list is in the `bp.ini` file.

service

A program on a Windows server system that runs in the background and performs some task (for example, starting other programs when they are needed). Services are generally referred to as daemons on UNIX systems.

session

An instance of NetBackup checking its schedules for backups that are due, adding them to its worklist, and attempting to complete all jobs in the worklist. For user backups and archives, a session usually consists of a single backup or archive.

Session (Vault)

A vault session consists of executing a particular profile or profiles.

shared drives

See “Shared Storage Option (SSO).”

Shared Storage Option (SSO)

A separately priced VERITAS software option that allows tape drives (standalone or in a robotic library) to be dynamically shared among multiple NetBackup and Storage Migrator servers.

This option is supported only on NetBackup DataCenter servers.

SMDR

Storage management data requestor, a Novell NetWare program that provides its services transparently to all SMS modules and lets remote and local modules communicate with one another.

SMS

Novell NetWare storage management services.

source volume group

A volume group from which Vault can select backups to duplicate. This parameter is used to restrict the list of backups from all backups that reside on media in any volume group to backups that reside on media in a single volume group. Where a volume group corresponds to a particular robot, the profile will duplicate only backups on media in that robot. The Source Volume Group is normally only specified if you have multiple robots attached to the same server, for example you want to duplicate backups that reside in robot 0 to media that reside in robot 1.

SSO

See “Shared Storage Option (SSO).”

stacker

Usually a small robotic library that contains one drive only. See “robotic library.”



standalone

A qualifier used with drives and media to indicate they are not associated with a robot. For example, a standalone tape drive is one where you must manually find and insert tapes before using them. A standalone volume is one that is located in a standalone drive or is stored outside of a drive and designated as standalone in the volume configuration.

status code

A numerical code, usually accompanied by a troubleshooting message, that indicates the outcome of an operation.

storage migrator

Refers to the VERITAS Storage Migrator line of hierarchical storage management products for UNIX and Windows. These products make extra room on a disk by transparently moving data to other storage and then transparently retrieving the data when it is needed by a user or application.

Storage Migrator is available only for NetBackup DataCenter servers.

storage unit

Refers to a storage device where NetBackup or Storage Migrator stores files. It can be a set of drives in a robot or consist of one or more single tape drives that connect to the same host.

SUSPENDED media state

If a volume is SUSPENDED, NetBackup can restore from it but cannot use it for backups. NetBackup retains a record of the media ID until the last backup image on the volume expires.

symbolic link

On a UNIX system, this is a pointer to the name of the file that has the source data.

TapeAlert

Allows reactive cleaning for most drive types and is a function of the tape drive.

tape format

The format that an application uses to write data on a tape.

tape marks

A mark that is recorded between backup images on a tape.

tape overhead

The space required for data that is not part of the backup images. For example, tape marks and catalogs of what are on the tape are considered overhead.

tape spanning

Using more than one tape to store a single backup image.

tar

Tape Archive program that NetBackup uses to extract backup images during a restore.

target

See “target service.”

target service

A Novell NetWare service that needs storage management. The SMS views all services (for example, print services, communication services, workstations) as targets.

Target Service Agent

A Target-service agent is a Novell NetWare agent that prepares the target's data for SMS during a backup and for the target during a restore.

TLD - Tape Library DLT

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

TLH - Tape Library Half-inch

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

This robot type is supported only by NetBackup DataCenter servers.

TLM - Tape Library Multimedia

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

This robot type is supported only by NetBackup DataCenter servers.

TL4 - Tape Library 4MM

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.



TL8 - Tape Library 8MM

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

timeout period

The period of time that an application has allotted for an event to occur.

TIR

See “true image restore.”

tpconfig

A Media Manager administration utility for configuring devices which is started from the command line. On UNIX, it has a character-based menu interface that can be run from terminals that do not have X Windows capabilities. tpconfig also has a command line interface.

transfer rate

The rate at which computer information is transferred between a source and a destination.

transport

See “robotic arm.”

true image restore

Restores the contents of a directory to what it was at the time of any scheduled full or incremental backup. Previously deleted files are ignored.

TS8 - Tape Stacker 8MM

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

TSA

See “Target Service Agent.”

TSD - Tape Stacker DLT

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

TSH - Tape Stacker Half-inch

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web site.

This robot type is supported only by NetBackup DataCenter servers.

unassigned media

Media that contain no valid images. A piece of unassigned media has an entry in the volumes database but no entries in the images database. Unassigned Media do not have a “time assigned” in the Media section of the GUI.

user interface

The program used to perform user backups, archives, and restores.

user operation

A backup, archive, or restore that is started by a person on a client system.

Vault

Vault is a separately-priced NetBackup option that provides offsite backup management. Vault automatically duplicates specified backup images, and automates the process of offsite media rotation (a critical component of any backup or disaster recovery strategy). Vault manages offsite storage and retrieval of media for original backups, duplicate backups, and catalog backups. Additionally, NetBackup Vault generates reports to track the location and content of each piece of media.

vault

In the context of the NetBackup Vault, a vault is logical entity associated with a particular robot that acts as a designated holding place for backups that will eventually be sent to a physical offsite vault. The term ‘vault’ is used to refer both to the process, and to the physical storage location of a set of tapes offsite.

vault process

Vaulting is the process of choosing backup images to duplicate or eject, optionally duplicating backups, ejecting duplicate or original media, storing it at an offsite location, and later returning expired media to your robot. Vaulting is an integral part of the disaster recovery process.

verbose flag

Configuration file entry that causes a higher level of detail to be written in the logs.



verify

An operation that compares the list of files that are actually on a volume with what NetBackup has recorded as being on it. The data that is on the media is not verified.

vmadm

A Media Manager administrator utility for managing volumes. It runs on UNIX and has a character-based, menu interface that can be run from terminals.

vm.conf

A Media Manager configuration file with entries that include the servers that can manage local devices and default media ID prefixes for media that do not contain barcodes.

volume

Media Manager volumes are logical units of data storage or cleaning capability on media that have been assigned media IDs and other attributes, which are recorded in the Media Manager volume database.

volume configuration

Refers to configuration information that is stored in the Media Manager volume database.

volume database

An internal database where Media Manager keeps information about volumes. All hosts (where Media Manager is installed) have a volume database. However, the database is empty unless the host is designated as a volume database host.

volume database host

The host (where Media Manager is installed) that contains information about the volumes that Media Manager uses in a device. Because NetBackup BusinessServer supports only a single server, the volume database host is always on the same server.

volume group

A set of volumes that are configured within Media Manager to reside at the same physical location (for example, in a specific robot).

volume pool

A set of volumes that are configured within Media Manager to be used by a single application and are protected from access by other applications and users.

wakeup interval

The time interval at which NetBackup checks for backups that are due.



wildcard characters

A character that can be used to represent other characters in searches.

Microsoft Windows

(noun) Describes a line of operating systems developed by Microsoft, Inc.

For more information on the Windows operating systems that NetBackup supports, refer to the VERITAS support web site at <http://www.support.veritas.com>.

Windows

(adjective) Used to describe a specific product or clarify a term. Some examples are: Windows 95, Windows 98, Windows NT, Windows 2000, Windows servers, Windows clients, Windows platforms, Windows hosts, and Windows GUI.

Windows servers

A term that defines the Windows server platforms that NetBackup supports; those platforms are: Windows NT and 2000.

Windows clients

A term that defines the Windows client platforms that NetBackup supports; those platforms are: Windows 95, 98, ME, NT, 2000, XP (for 32- and 64-bit versions), and LE.

Windows Display Console

A NetBackup-Java interface program that runs on Windows 2000, NT, 98, and 95 computers. Users can start this interface on their local system, connect to a UNIX system that has the NetBackup-Java software installed, and then perform any user operations that their permissions allow.

WORM media

Write-once, read-many media for optical disks. NetBackup Business Server does not support WORM media.

xbp

The X Windows-based backup, archive, and restore program for users on NetBackup UNIX clients.





Index

A

- Activity Monitor
 - Vault support 103
- Add
 - alternate media server names 95
 - license key 11
- Addresses for email notification 92
- Advanced configuration 76
- Alternate names
 - add 95
- Architectural services 188

B

- Backup Images, sending to Volume Pools 31
- Best Practices 25
- bpdbjobs
 - changes for Vault 134
- bpexpdate 77
- bpvault
 - upgrade to NetBackup Vault 159

C

- Catalog Backup
 - ensuring media for 106
- Catalog Backup tab 64
- Catalog Backups
 - large 67
- Catalog node
 - Inline Tape Copy 81
- Cautions 13
- Choose Backups tab 52
- Collect NetBackup information 16
- Commands
 - bpadm 45
 - bpbackupdb 39
 - bplabel 107
 - bpstulist 19
 - get_license_key 10
 - Print 129

- query acs all 20
- vltadm 45
- vltrun 88
- vmquery 107

Complete Inventory List for Vault 113

Configuration

- Catalog Backup tab 64
- Choose Backups tab 52
- Duplication tab 54
- Eject tab 68
- Media Manager 20
- methods 45
- network 20
- profile 21, 51
- Reports tab 69
- robots 46
- using vltadm 45
- vault 21

Continue

- for multiple copies 74

Copy Profile 94

Create

- profile 50
- vault 47

D

- Data
 - location 166
- Debug logs 139
- Detailed Distribution List for Vault 110
- Dialogs and screens
 - Advanced Configuration 60
 - Alternate media server names 96
 - Catalog Backup 64
 - Change Report Titles 71
 - Choose Backups 52
 - Duplication 56
 - Eject 68
 - Email 92



-
- Multiple Copies 59
 - New Profile 50
 - New Vault 47
 - New Vault Robot 46
 - Reports 70
 - Directory structure 181
 - Distribution List for Robot 111
 - Distribution List for Vault 110
 - Duplication
 - Advanced Configuration dialog 59
 - basic 56
 - Multiple Copies dialog 58
 - multiple duplication rules 62
 - multiplexed 82
 - rules 55
 - subscreens 57
 - through Catalog node 83
 - where possible 82
 - Duplication tab 54
- E**
- Edit Vault or Profile 93
 - Eject tab 68
 - Email notification 92
 - Error codes 135
 - Error codes, extended 105
 - EXIT status 135
 - Expiration 58
 - Expiration date 58
 - Extended error codes 105
- F**
- Fail
 - for multiple copies 74
 - Full Inventory List for Vault 112
 - Functional Design
 - client/server 191
 - services interactions 190
 - Functional design
 - architectural services 188
 - overview 187
 - technical components 194, 195
 - technical design issues 194
- H**
- Help
 - vltadm 128
- I**
- Inline Tape Copy
 - Catalog node 81
 - overview 73
 - policy node 79
 - Install
 - on UNIX systems 10
 - on Windows systems 11
 - prerequisites 10
 - supported platforms 9
 - supported robots 9
 - Inventory List for Vault 112
- L**
- License key
 - add 11
 - remove 12
 - Load Balancing 39
 - Log Files
 - set duration 99
 - Log files 96
- M**
- Media
 - for catalog backup 106
 - Media Manager
 - configuration 20
 - Media Server Names
 - add alternate 95
 - Menu User Interface
 - accessing 118
 - bpdbjobs 134
 - catalog backup 124
 - changes in vmadm 131
 - choose backups 122
 - duplication 123
 - duplication items setup 123
 - eject 125
 - help for vltadm 128
 - overview 117
 - profile management 121
 - reports 126
 - robot management 120
 - vault administration 118
 - vault management 120
 - vault properties 119
 - vltadm 118
 - vltopmenu 129
 - volume pools 125
 - Move vault to different robot 94
 - Multiple Copies
 - dialog 58
 - select continue or fail 74

- Multiple copies
 - advanced configuration 76
 - creating outside Vault 79
 - overview 73
- Multiple media types, managing 15
- Multi-Robot Configurations
 - good 37

N

- Names
 - Alternate Media Server 95
- Network configuration 20
- Notes
 - Activity Monitor 103
 - alternate media server 95
 - alternate media server names 60
 - automatic report mode 72
 - configure printers 21
 - copy profile 94
 - expired backups 58
 - move a vault 94, 95
 - multiple copies 59
 - multiplexing 58
 - number of copies 63
 - prepare to use email 92
 - suspension 69
 - unassigned media 65
 - Vault session and MUI 88
 - volume pool, selecting 66
- Notification, email 92
- Notify Scripts
 - for specific Profile 102
 - for specific Robot 102
 - for specific Vault 102
 - order of execution 103
 - using 100

O

- Off-site volume group 48
- Order of execution for scripts 103

P

- Picking List for Robot 110
- Picking List for Vault 111
- Planning 15
- Platforms supported 9
- Policy node
 - Inline Tape Copy 79
- Policy type 17
- Preview Vault Session 88

- Primary copy 58
- Profile
 - configure 51
 - copy 94
 - create 50
 - edit 93
 - Notify Script for 102
 - overlap time window 27
 - print information 93

R

- Recovery
 - freeze duplicate copy 176
 - identify damaged media 174
 - request media return 176
 - steps 173
- Remove
 - license key 12
- Reports
 - consolidate 113
 - detailed 111
 - distribution 115
 - for media coming on site 111
 - for media going off site 110
 - recovery report 113
 - reprint 115
 - run from command line 114
 - types 109
- Reports tab 69
- Resume Vault Session 91
- Retention level 57
- Return
 - request return of media 176
- Revault media 178
- Robot
 - configure for Vault 46
 - Notify Script for 102

S

- Scheduler
 - run Vault Session 91
- Services interactions 190
- Session, Vault
 - definition 87
 - locking 138
 - preview 88
 - resume 91
 - run 88
 - run from GUI 89
 - run from scheduler 91



- stop 105
- Slot ID 48
- Source volume group 53
- Status Codes 135
- Summary Distribution List for Vault 110
- Supported platforms 9

T

- Time Window, overlapping 27
- Troubleshooting
 - duplicate media 136
 - eject issues 138
 - ejecting tapes in use 138
 - error codes 135
 - logs 139
 - media missing in robot 135
 - no duplicate progress message 137
 - overview 135
 - robot goes offline 137
 - session lock 138

U

- Unfreeze media 178
- Uninstall 12
- UNIX systems
 - install Vault on 10
 - uninstall Vault 12
- Upgrade
 - feature comparison 160
 - location of files 166
 - overview 159
- Uses for Vault 1

V

- Vault
 - accessing 3
 - create 47
 - edit 93
 - fit with NetBackup 3

- move to different robot 94
- Notify Script for 102
- overview 2
- print information 93
- Vault Session
 - definition 87
 - lock 138
 - preview 88
 - resume 91
 - run 88
 - run from GUI 89
 - run from scheduler 91
 - stop 105
- vlt_ejectlist_notify 101
- vlt_end_notify 101
- vlt_endeject_notify 101
- vlt_start_notify 101
- vltadm 118, 142
 - when to use 45
- vlteject 144
- vlthinject 148
- vltoffsitemedia 150
- vltopmenu 129, 153
- vltrun 154
- vmadm
 - changes for Vault 131
 - special actions menu 131
 - volume configuration 131
- Volume Pools, naming 34

W

- Windows system
 - install Vault on 11
- Windows systems
 - uninstall Vault 12
- Working files
 - set duration 99